An Interop Publication

CONNEXIONS®

The Interoperability Report

August 1991 Volume 5, No. 8

ConneXions — The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

ConneXions is published monthly by Interop, Inc., 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. 415-941-3399. Fax: 415-949-1779. Toll-free: 1-800-INTEROP.

ISSN 0894–5926

From the Editor

Our long-running series *Components of OSI* continues this month with a look at *ODA,* the Office/Open Document Architecture. ODA is an international standard that provides for open interchange of revisable and formatted documents, and is defined jointly by ISO and CCITT. As Jon Stewart explains, ODA represents yet another important piece of the OSI architecture.

Gary Malkin gives a brief introduction to TCP/IP naming and addressing. In a future edition we will tackle the more complex issue of how to expand the existing IP address space, as well as look at the OSI NSAP address assignment problem.

ConneXions has featured network "profiles" since we first started publication in 1987. This month, we've included two such profiles. First, Tracy LaQuey Parker gives an outline of the *Texas Higher Education Network* (THEnet) and the related TENET project. In light of last month's special issue on the Changing Face of the Internet, it is interesting to note that TENET plans to give Internet access to over 6000 public schools. The Internet is indeed changing.

Our second network profile comes from a very different part of the world. Mike Lawrie of Rhodes University in South Africa describes efforts to build a national research and academic network in the face of political and regulatory obstacles, both national and international.

As usual, we include a number of brief announcements, call for papers and the like. In particular, I'd like to draw your attention to the note on page 25 about the new TCP/IP CD ROM which is available from SRI International. Imagine having all the RFCs and "stacks" of other relevant documents available on your desktop computer system without filling up your hard disk! Personally, I don't yet have the technology to read the CD, but I have asked someone with the appropriate hardware to send me a "CD Review" for a future issue. (I tried playing the CD on my stereo, but it sounds terrible! :–)

INTEROP 91 is only two months away. In July, the shownet engineering crew assembled at the San Jose Convention Center for the annual INTEROP "Wiring Party." Meanwhile, participants for the various *Solutions Showcase Demonstrations* are hard at work preparing for this year's demos. Descriptions of the demonstrations can be found on page 28.

Last month, we brought you an adapted chapter from Carl Malamud's soon-to-be-published book, *STACKS—Interoperability in Today's Computer Networks.* We neglected to tell you that the book will be published by Prentice-Hall, and that the chapter was printed with permission. Sorry, Paul.

Components of OSI: Document Interchange Using ODA

by Jon A. Stewart, Digital Equipment Corporation

Introduction

The *Office/Open Document Architecture* (ODA) is an international standard that provides for open interchange of revisable and formatted (final form) documents. ODA is defined jointly by ISO (ISO 8613) and CCITT (T.410 series) for use at the application level of OSI. The text of ISO 8613-1 to ISO 8613-8 is identical (when applicable and except for stylistic conventions) to corresponding CCITT Recommendations T.411 to T.418.

ODA was initially planned in 1981 for interchange of office documents typical of word processing at that time. However, the rapid progress of desktop and electronic publishing so influenced ODA development that, as published in 1989, ODA can describe complex technical documents mixing character and graphics text. *Open Document Architecture* is the name used in the CCITT T.410 series and it has now been approved by ISO in order to reflect a much wider scope for ODA applications.

The ODA standard provides for two data interchange formats to be communicated by network transfer or storage media exchange:

- *Office/Open Document Interchange Format* (ODIF)—an application of ASN.1
- *Office/Open Document Language* (ODL)—an application of the Standard Generalized Markup Language (SGML) and the SGML Data Interchange Format (SDIF, an application of ASN.1)

Only ODIF is specified by the CCITT T.410 series; and, to this date, only ODIF is being utilized in the released products for transparent ODA converters via X.400 gateways.

Text and information representation

The ODA standard defines "document" as "A structured amount of information intended for human perception, that can be interchanged as a unit between users and/or systems." ODA conveys such information in a variety of coded *content elements*—such as character text, geometric graphics, raster images—all of which can be mixed within the same document. Because of the flexibility provided, typical ODA documents are interchanged with *processable* or *formatted processable content elements* (PCEs, FPCEs.) After interchange, further manipulation of this content is possible—such as editing, spell checking, content-based retrieval, scaling, clipping, changing of fonts and other graphic renditions, reformatting, etc. In ODA layout, two steps are involved when the PCEs are to be presented/imaged:

- First, the content layout process is performed to size and position the content in the layout object allocated by the document layout process; this step may derive an FCE or FPCE from a PCE, or it may operate directly on an FCE or FPCE already available
- Then the laid out content is presented by the imaging process using presentation attributes and control functions (that may have been inserted in the first step).

Note that, for character text, the first step requires font metrics (for precise sizing and positioning, as when kerning) but character/glyph images are *not* required until imaging is performed. Thus it is possible to interchange a formatted document without having the font available, indicating in the document the precise font requirements for imaging in terms of the font standard (ISO 9541).

Processes such as hyphenation (dependent on language) and imaging (dependent on specific devices) are so numerous and variable as to be extremely difficult to standardize. The ODA approach to solving this open interchange problem was to allow for the retention of layout decisions made by non-standardized processes not available to the system receiving the document. Character text FPCEs carry the formatting decisions (inserted as various control codes and escape sequences) but at the same time mark them for easy removal. The FPCE can thus be imaged as intended (if the receiver can interpret the control functions); but, when required, the PCE can be recovered for editing and other processing, including reformatting.

To allow for addition of new content/information types, ODA was designed to provide a generalized interface to content that allows overall document description (macro structure) to proceed independently of content-specific (micro structure) description. Each "content architecture" must define its micro structure as follows:

- The PCE (and FCE and FPCE when applicable) data structures
- The presentation attributes that control the sizing and positioning of content in layout objects
- The presentation attributes (and control functions) that control the imaging process
- Rules (if applicable) for recovering PCE text from FPCE text
- Guidelines for implementing the content layout process (for example, a line layout process including kerning for character text; coordinate system relationships—e.g., CGM includes one for its pictures that must be related to the ODA layout coordinate system)

ODA currently provides three content architectures:

- Character text (based on ISO character set standards and invocation rules defined in ISO 2022, and the ISO 9541 font standard)—PCE, FCE and FPCE are available
- Raster graphics based on CCITT facsimile recommendations T.4 and T.6—defines an FCE and FPCE form (scaling, clipping permitted)
- Geometric graphics based on the binary encoding of the *Computer Graphics Metafile*, ISO 8632—only the FPCE form is available (scaling, clipping, rotation permitted)

Document structure

ODA provides the data models and structural rules for organizing documents into both *logical objects* (such as chapters) and *layout objects* (such as pages). The objective of such dual structuring is to provide for interchange that allows both revisability and presentation (imaging) as intended by the originator. Accordingly, a document may be interchanged in any of three different states:

- Processable—contains a *specific logical structure* and associated PCEs (and generic layout structure and styles if formattability is desired)
- Formatted—contains a *specific layout structure* and FCEs (allocated to layout objects in that structure)
- Formatted processable—contains both structures and FPCEs referenced from both (see Figure 1)

continued on next page

Document Interchange Using ODA *(continued)*

Example Formatted Processable Document

Figure 1 shows an example of a formatted processable ODA document. This highly contrived example was designed to be small and simple but illustrate many of the features of ODA. The FPCEs shown are associated with both the logical object named in the indented outline at right *and* with the layout object (the rectangular boundary) in which it is imaged in the figure.

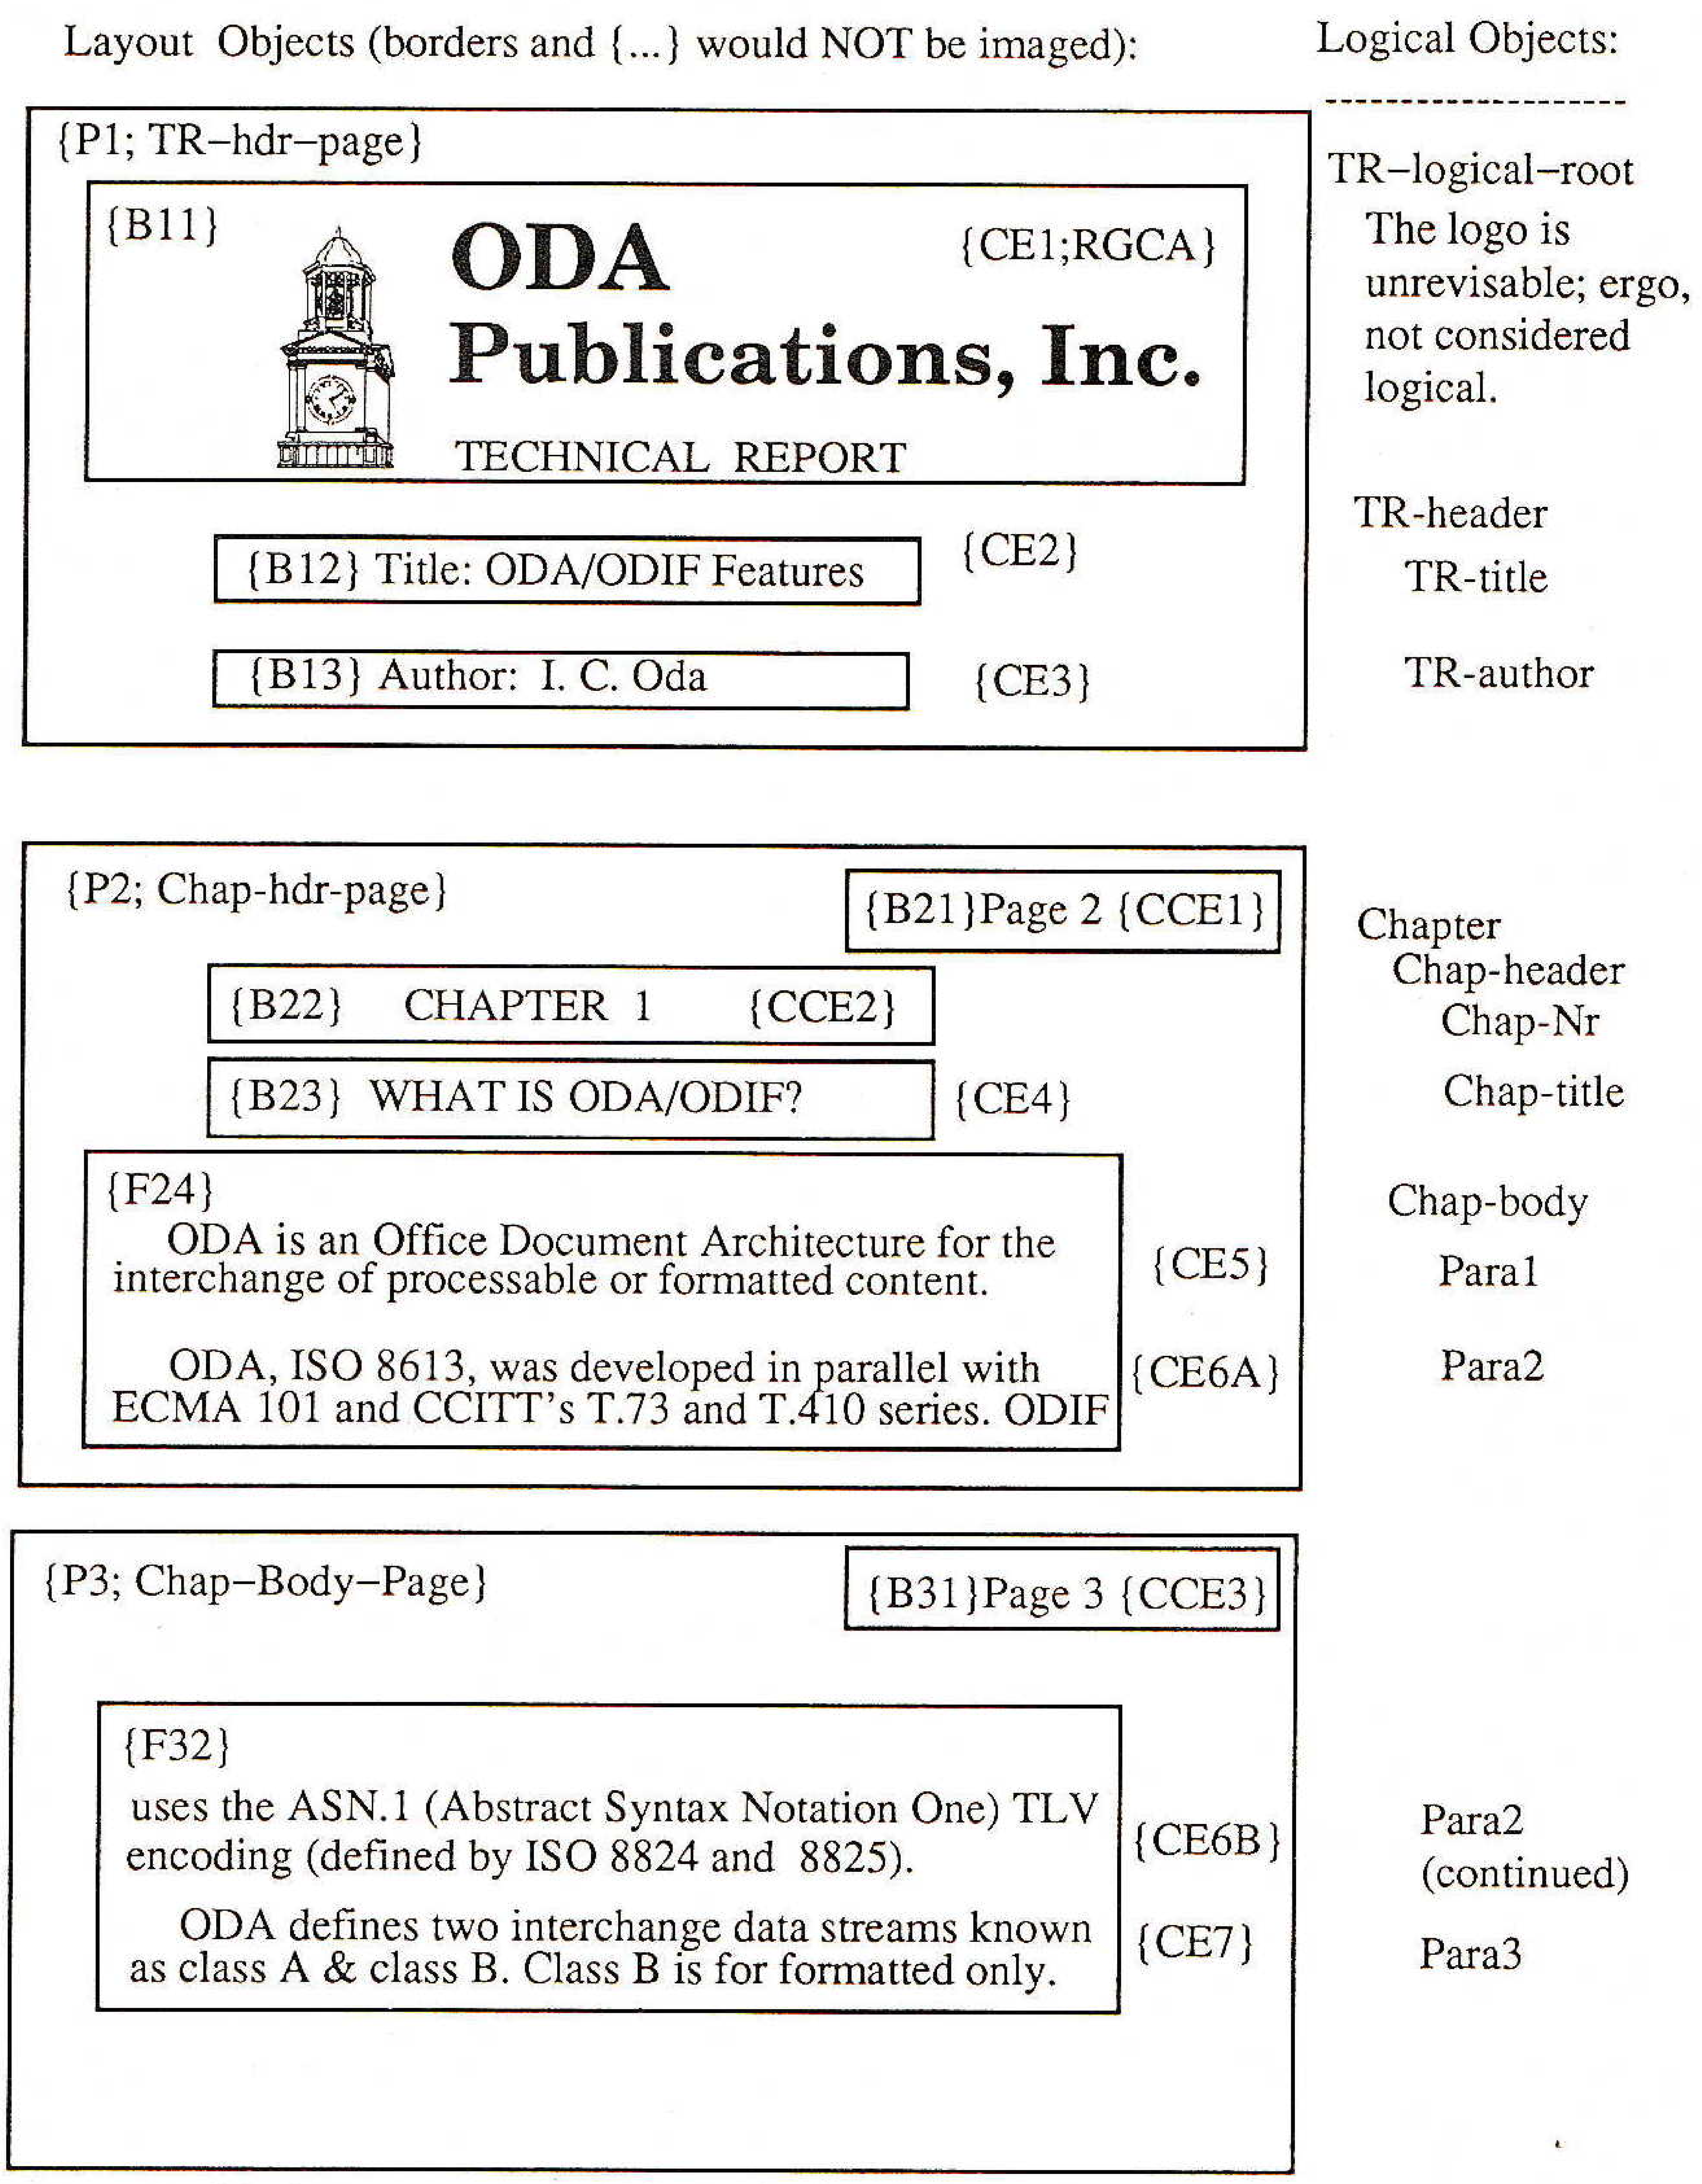

Notation: {...} indicates comments NOT to be imaged.
P=Page, F=Frame, B=Block, CE=Content Element, CCE=Computed CE
RGCA=Raster Graphics Content Architecture (T.4/T.6 fax image)
xx;object-type/class (e.g., P1;TR-hdr-page = page P1 of class indicated)

Figure 1: Example Formatted Processable (FP) Document

Concurrent Structure

Figure 2 clearly shows that there is dual (concurrent) structure within the document:

Non-terminal nodes of the trees in Figure 2 represent *composite objects* and both trees terminate on *basic objects* that reference the same content elements. The layout tree is rooted at a special composite object called the *document layout root* and the logical tree has a *document logical root.* Content can only be associated with basic objects, and one basic object can reference multiple content elements of the same content architecture (as is the case with logical object *Para2* and layout objects *F24* and *F32*).

The Figure 1 formatted processable (FP) document has been completely laid out for imaging on three classes of page (*TR-hdr-page, Chap-hdr-page, Chap-body-page*). Content elements (all FPCEs) have been prepared to fit into the blocks according to the rules of each content architecture, including insertion of control functions to provide positioning, font selection, and other renditions. One paragraph (*Para2*) wouldn't fit completely in *F24* of *Chap-hdr-page* and had to be split into two content elements *CE6A* and *CE6B*. *CE6A* appears in block *F24* of the *Chap-hdr-page* and the remainder of *Para2, CE6B,* has been placed into *F32* of a *Chap-body-page*.

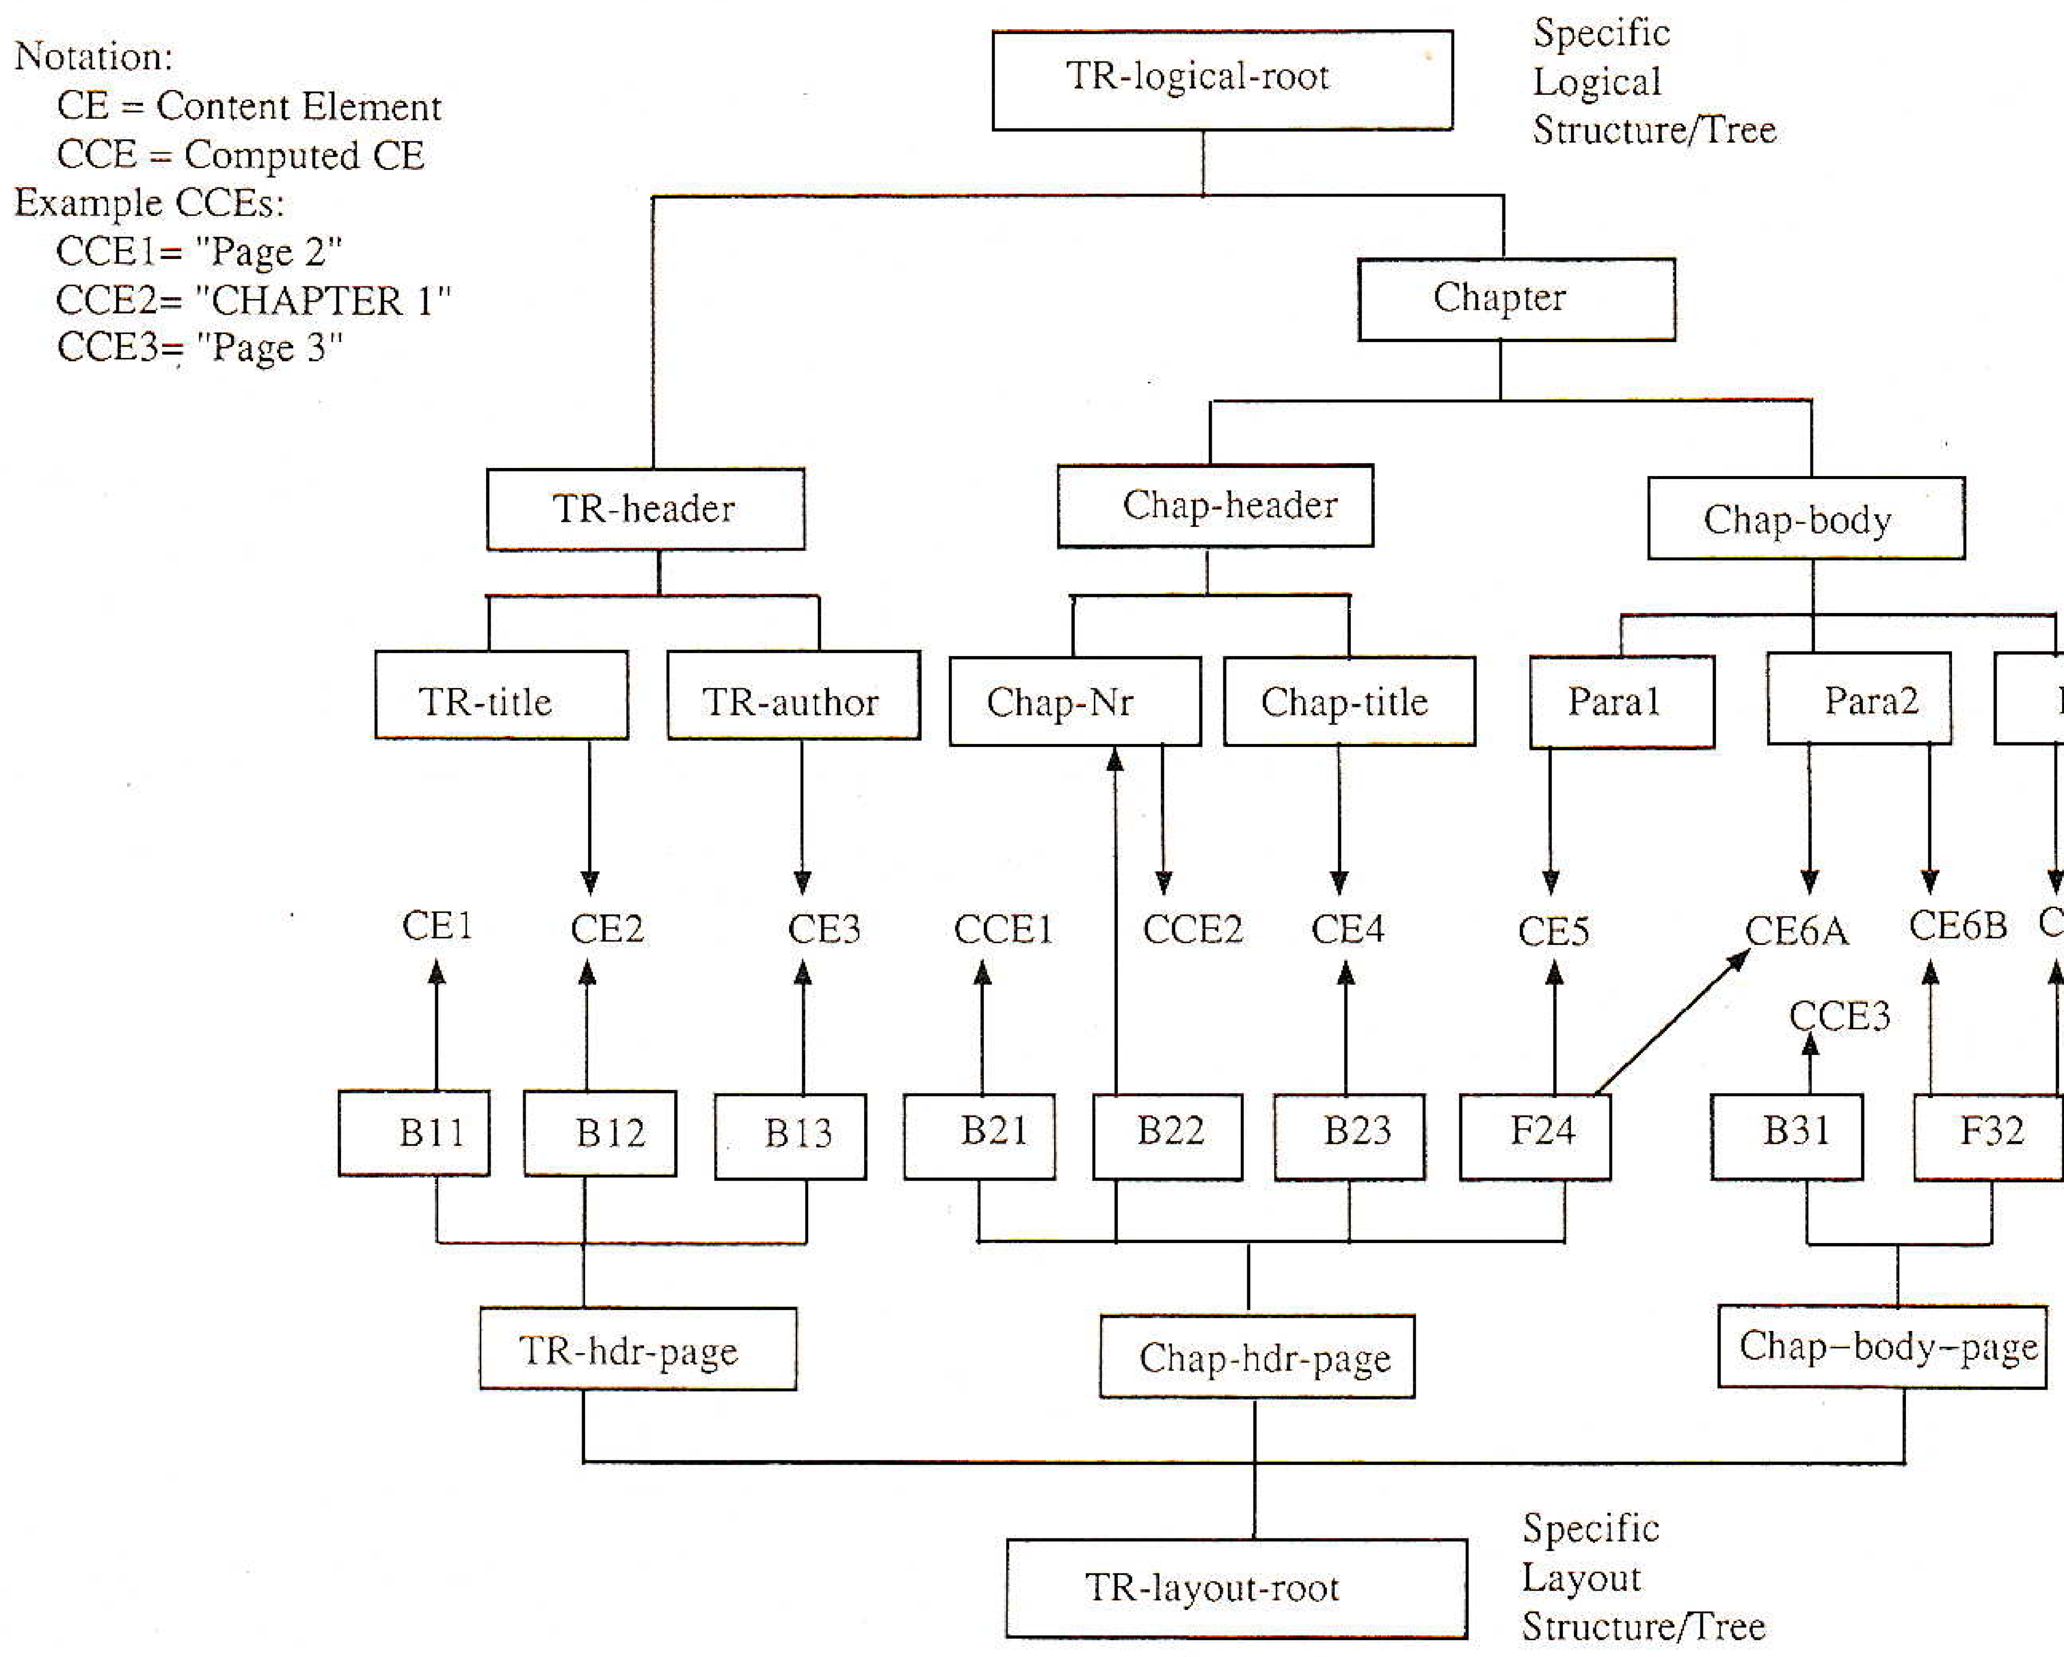

Figure 2: Logical and Layout Trees for the Figure 1 Document

Computed/generated content

Content elements generated by evaluation of numeric and string expressions defined in ODA have many uses in document processing. Although processes for evaluation of the expressions are *not* defined in the standard, the intent for evaluation is expressed in the interchange format, and may be utilized by implementations to produce the desired result—for example, automatic chapter or page numbering. Figure 2 shows three such *computed content elements* (CCEs).

Chapter number is evaluated for this document by assigning a *binding name* (*Chap-Nr*) and using a numeric function INCREMENT PRECEDING to evaluate the *binding value* of the current instance of *Chap-Nr.* In this case, the PRECEDING function searches in reverse sequential order back through the specific logical structure until an object is found that has a binding for *Chap-Nr,* and when it is found the binding value for that is incremented to produce the binding value for the current instance of *Chap-Nr.* If the document root is reached without finding an occurrence of the binding name (as it would be for chapter 1 in this example) then a value of 0 or "null" (if in a string expression) is used. Evaluation of *Chap-Nr* would be expected whenever a process (probably an editor) either deleted or added a chapter to the document.

Page-Nr is treated in a similar fashion to *Chap-Nr,* with the difference that the tree searched is that of the specific layout structure. Evaluation of *Page-Nr* would be expected to be performed by the ODA layout process.

continued on next page

Document Interchange Using ODA *(continued)*

Generic logical structure

A diagrammatic notation for generic structure is used in the ODA standard. Figures 3 and 4 use this notation to show the generic structures of the Figure 1 document. This notation is related to the ASN.1 encoding rules that provide for structured data elements as follows:

REP	=	REPetition (1 or more instances) of data elements
CHO	=	CHOice from enumerated list of data elements
OPT	=	OPTional data element
AGG	=	set of data elements (AGGregate, any order)
SEQ	=	SEQuence (required order) of data elements
OPT, REP	=	OPTional REPitition (0, 1 or more elements)

The Figure 1 document is a specific document instance that can be generated from the generic logical structure provided in Figure 3. To generate the logical structure of this document, start at the document logical root and follow the nodes down and to the left using the structured data rule appearing above the connecting lines. When a terminal (basic) node is reached by following left hand branches, return to the composite object above it and repeat for the next subordinate of that composite object, in left to right order at that level. This produces:

TR-logical-root

TR-header, TR-title, TR-author

[repeat 1] *Chapter*

Chap-header, Chap-Nr, Chap-title

Chap-body [repeat 3] *Para1, Para2, Para3*

Purpose

Generic logical structure is provided for by ODA so that appropriately designed editing processes can control structural and content modification in a consistent way. Generic logical structure serves the same purpose as that of a language grammar defining syntactically valid source programs that may be written in a programming language. Most compilers won't allow a programmer to invent a new statement form or place a statement in an arbitrary position within the source. Neither should a well designed editing process allow a document author/preparer to violate the constraints of structure imposed by the document designer in the form of generic logical structure.

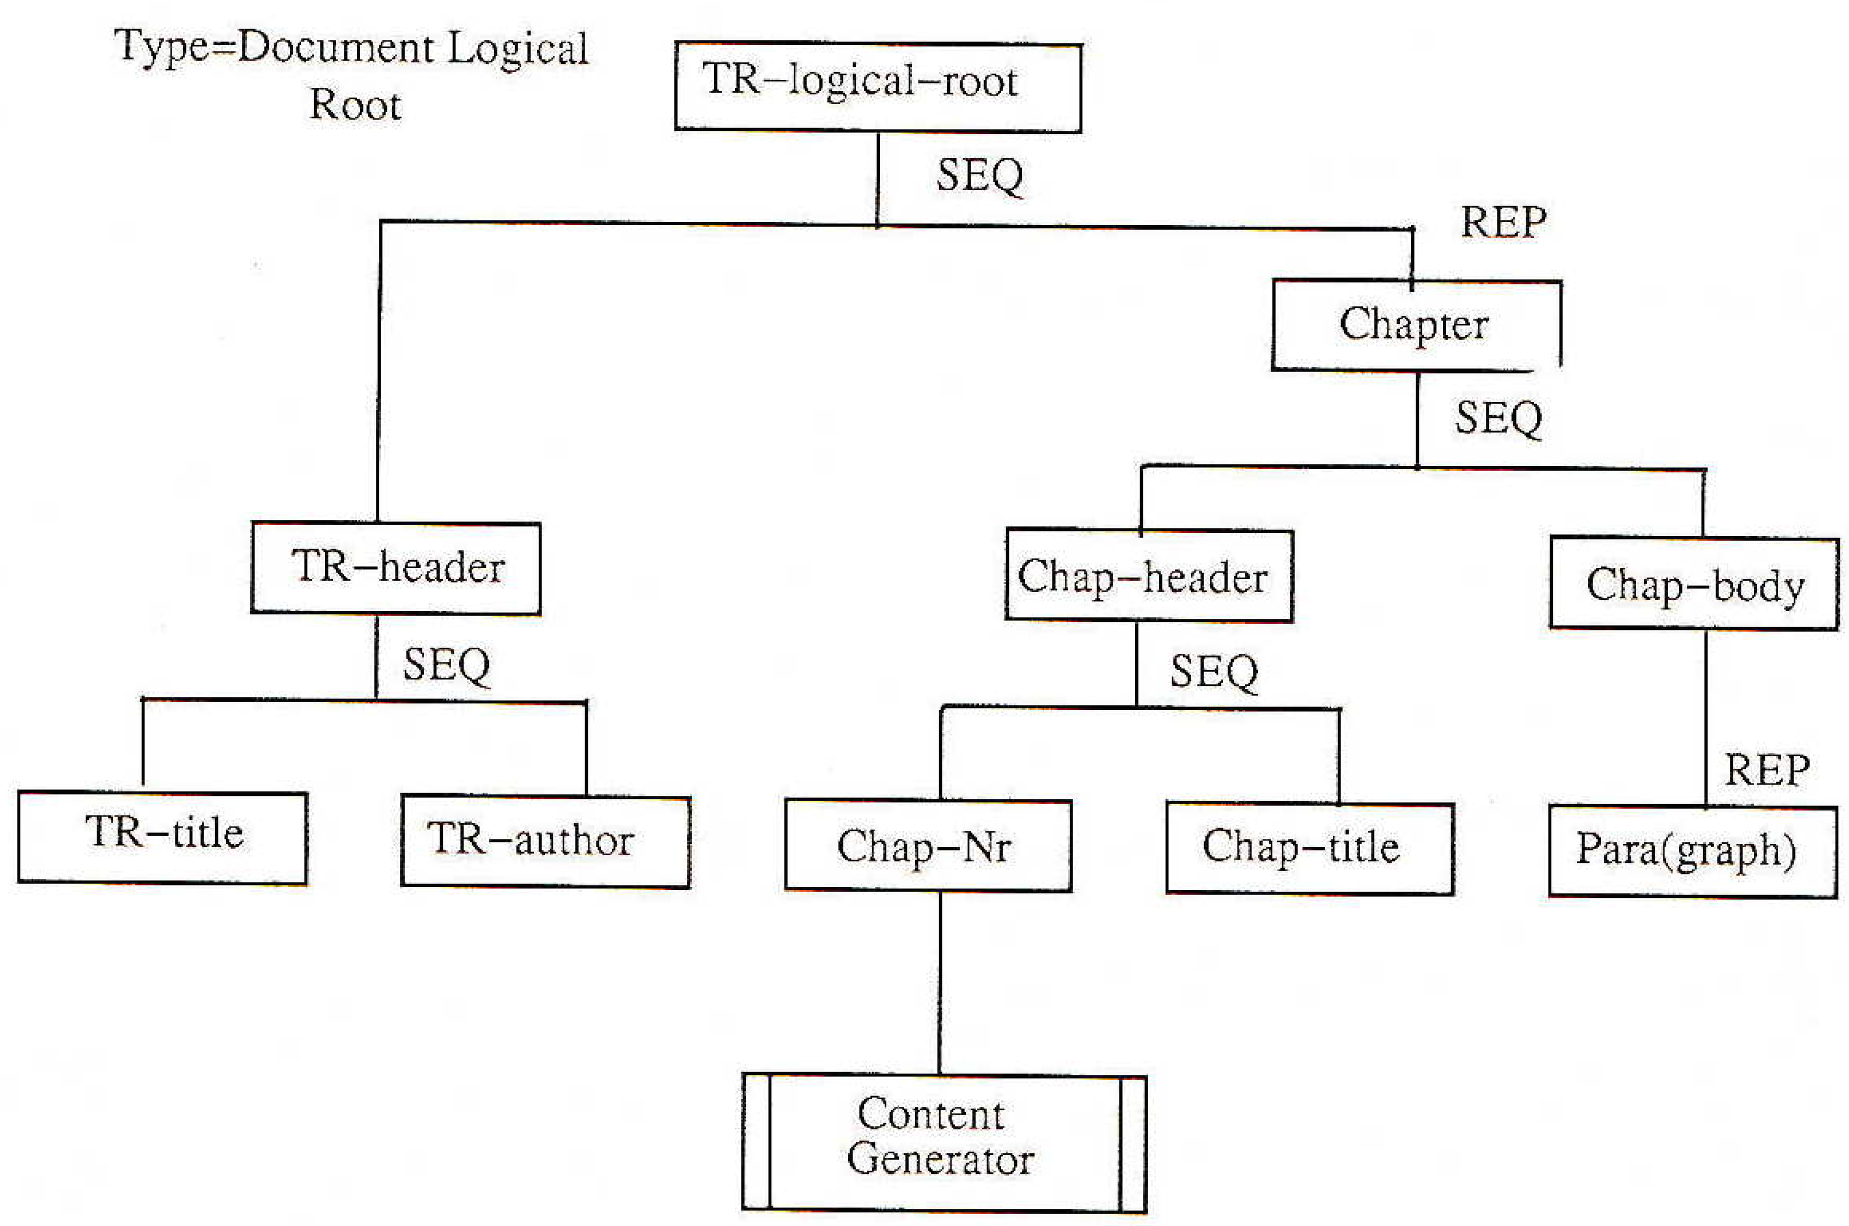

Figure 3: Generic Logical Structure for the Figure 1 Document

Let's say the author wants to insert a new chapter into the existing Figure 1 document. Table 1 below illustrates two points at which a Chapter insert would *not* be permitted and one point at which it would.

Chapter Insert:	*Logical Objects of Figure 1 Document*			
	TR-logical-root			
1) NOT ALLOWED————>				
		TR-header		
			TR-title	
			TR-author	
		Chapter		
2) NOT ALLOWED————>				
			Chap-header	
				Chap-Nr {binding value 1}
				Chap-title
			Chap-body	
				Para1
				Para2
				Para3
3) ALLOWED————>		*Chapter*		
			Chap-header	
				Chap-Nr {binding value 2}
				Chap-title
			Chap-body	
				Para1-2
				...

Table 1: Use of Generic Logical Structure

Logical semantics

ODA actually defines only two types of logical object—*composite* and *basic*. An application of ODA is free to define the meaning (to its application-specific processes) of the logical objects it builds up from these structural components. Names such as "paragraph" or "chapter" have no meaning in ODA itself—that is, they do *not* have standardized sets of attributes ("logical semantics"). In the ODA standard each logical object is simply identified by an "object identifier" and its type (basic or composite, generic or specific) determines the attributes permitted—rather than specifying the values of such attributes, as would be the case for a standardized "paragraph" or "chapter" object.

DAP

Implementors of ODA systems are cooperating through workshops such as the NIST ODA SIG (part of the *NIST Workshop for Implementors of OSI*) to standardize logical and layout semantics beyond that provided by the base ODA standard. This is being done by means of a specification provided for by the standard, the *Document Application Profile* (DAP). A DAP effectively constrains (selects values from the many permitted by the standard) attribute values of objects including references to subordinate objects. In this way the DAP effectively defines standardized objects (both layout and logical) that facilitate interoperability between processes (such as editors, formatters) that agree to conform to the DAP.

The Layout Processing Model

Although ODA does *not* standardize a process for layout (formatting) text it *does* precisely define its input and output and control parameters (the styles). The layout process requires input as follows:

- Specific logical structure and logical content (PCEs) associated with basic objects
- Layout object class descriptions (generic layout)

continued on next page

Document Interchange Using ODA *(continued)*

- Layout directives/styles to control creation of document-level layout objects (e.g., page-sets, pages, frames, blocks)
- Presentation attributes/styles to control content-level layout processes (and imaging)

The output of the ODA layout process is a formatted or formatted processable document (Figure 1) in which content elements have been sized and positioned for imaging and a specific layout structure has been created to "hold" the formatted content.

Note that a processable document that contains generic layout structure and layout and presentation styles (or references a "resource document" containing this information) meets the requirements for "formattability" in accordance with the guidelines provided for the ODA layout process. Most current ODA document interchange is of this processable/formattable variety. Figure 4 shows a generic layout structure that could be used with the Figure 1 document.

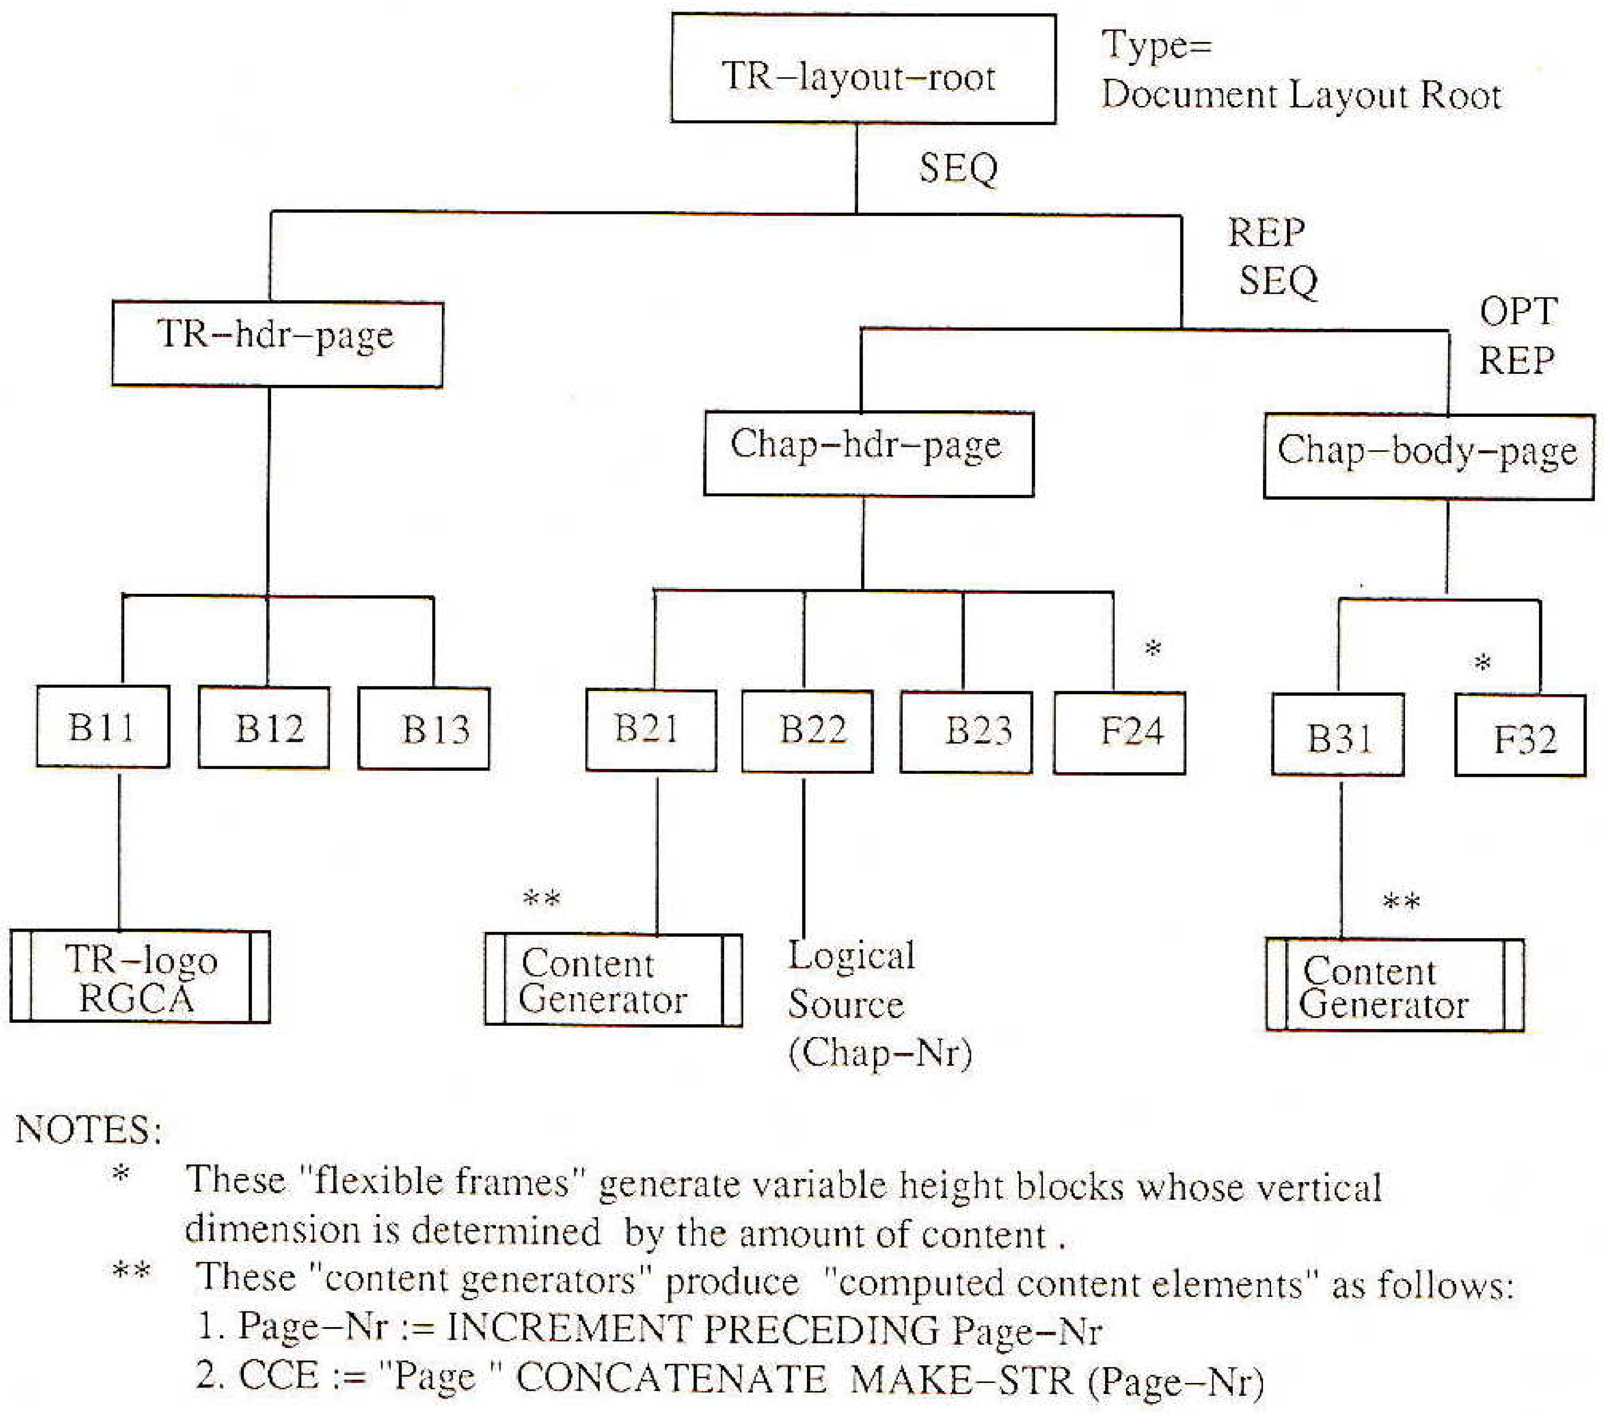

Figure 4: Generic Layout Structure for the Figure 1 Document

Layout and Presentation styles

The layout process is guided by the layout object class descriptions (generic layout) and layout/presentation attributes provided from logical and layout objects. Presentation attributes appear only on "basic objects" (i.e., those containing content that needs content layout and imaging). Layout directives may appear on composite and basic logical objects. All such attributes are typically collected into "styles" that are then referenced by object identifier from the appropriate objects.

"Layout styles" are used to control the overall document layout process (creation of specific layout structure) and "presentation styles" control the content layout processes (i.e., formatting and placing content into the allocated layout objects) and imaging. Logical objects may control their own document layout destination (and destiny) by using such layout directives as the following (small sample):

- Indivisibility—don't break the logical object across layout objects
- Layout object class—create a new layout object as identified to contain only this logical object
- Layout category—place the logical content in a frame that has "permitted category" to match

Table 2 below indicates the layout directives that the document designer of the Figure 1 document might use to control placement of the content (supplied and computed) elements into a specific layout structure controlled by the generic layout structure of Figure 3.

Logical Object	*Layout Directive*		
TR-header	Layout-Object-Class	=	*TR-hdr-page* *
Chap-header	Layout-Object-Class	=	*Chap-hdr-page*
Para(graph)	Layout-category	=	*Para-Frame* **

Notes:

* *TR-logo* is automatically laid out in a raster-graphics block when the layout object for *TR-header* is created.

** Frames allowed to receive *Para*(graphs) would specify a Permitted-Category of *Para-Frame*.

Table 2: Layout Directives for Figure 1 Document

General flow of the layout process

Before layout starts, any existing specific layout structure is deleted and any FPCE is restored to its PCE form—for example, by removing hyphens inserted by the previously executed content layout process and recombining previously split logical content, such as *Para2* in the example of Figure 1. The following offers a very general description of how the layout process then proceeds:

- *Current logical position* is initialized to *document logical root* in the specific logical structure
- *Current layout position* is initialized to *document layout root* in the generic layout structure
- The current logical object is processed
- Layout directives of the logical object are applied to create and select layout objects suitable for holding the content (e.g., processing the *TR-header* logical object would generate a *TR-header-page* layout object)

Note: Composite logical objects are examined for layout directives/styles that may cause creation of new layout objects and set up attributes (e.g., block offsets, interline spacing) that influence the processing of the basic objects that contain the content (requiring layout).

- When a basic logical object is processed the associated content is laid out according to the rules of the content architecture into the current layout object (if its permitted category allows), under control of the layout and presentation attributes in effect
- As layout objects are filled they are closed, and the generic layout structure and references to it through layout directives are used to establish new layout objects and current layout position
- Current logical position is advanced in sequential logical order until the last object is processed

Planned extensions to ODA

Many extensions of ODA (beyond the 1989 version) are essentially complete and in final stages of approval/publication. These include:

- A color amendment providing document level color (e.g., background color of a frame) and content element color (e.g., the color of the glyph) for high quality printing applications ("paper color") as well as video displays ("glass color")

continued on next page

Document Interchange Using ODA *(continued)*

- A tiled raster graphics content architecture that provides for splitting of very large raster graphics (RGCA) content into manageable arrays of "tiles" (sub pictures) of much smaller size than the complete picture
- Document security provisions allowing protection (by process, including encryption) of elements/objects of the ODA document at essentially any level within the document (e.g., a particular content element, a chapter, a page, etc.).

Much interest has been shown (and much preliminary work done) in extending ODA into application areas such as the following:

- Data interchange/sharing (including VTX, EDI, SQL Access)
- Business graphics and other "formatting" transforms
- Report generation, spreadsheeting, list processing (and associated tabular layout)
- Mathematical equations and chemical formulae
- Document processing (external references, editing, office procedures, hypertext, multimedia, data entry/forms processing)

The work of all but the last of the above items is proceeding under the collective title of "data in documents" in a Special Working Group (SWG-DiD) of ISO/IEC JTC1/SC18 WG3 and WG5 and is at the status of working draft review and debate. The last group of items is collectively referred to as *ODA Document Processing* and is proceeding in SWG-X of WG3. All currently approved extension work is in relatively early stages of development (if you're interested, join up!) and will probably take several years to complete.

Product support

The first ODA gateway/converter products are now reaching the market and continued progress is being made in various implementation efforts of the version 1 (1989) products. Some of these are in cooperative development efforts such as the newly formed *Open Document Architecture Consortium* (ODAC) whose members are Bull, DEC, IBM, ICL/Fujitsu, Siemens/Nixdorf and Unisys. XEROX and DEC have already released an ODA converter and X.400 gateway product handling documents conforming to the Q112 DAP. Several others have announced Q112 converter products; and, at least one Japanese product for Q111 (character text only) documents has been announced.

For more information about ODA, contact:

Jon A. Stewart, International Software/Standards Analyst
Digital Equipment Corporation, Mailstop MKO 1-2/F31
Continental Blvd.
Merrimack, NH 03054
Phone: +1 (603) 884-5967
Fax: +1 (603) 884-1036
E-mail: STEWART@CUERVO.DEC.COM

JON STEWART received his BS in Physics (Tulane University, 1956) and MA in Astronomy/Astrophysics (Indiana University, 1958). He programmed stellar model computations at IU as a graduate student and was also involved at IU in some of the earliest (1961) research in computational linguistics. He has been a computer science educator and research associate/consultant, a DEC software support and educational marketing specialist; and, for DEC since 1978, a software designer. Since 1985 he has been involved in establishing DEC's international software engineering architectures and standards, analyzing potential international software alliances, and representing DEC in the development of international standards committees for document interchange and processing. He is the Chair of the task group in ANSI X3V1 that coordinates the USA's participation in ODA.

What's in a name?

by Gary Malkin, FTP Software

Introduction

As the field of computer networking has evolved, it has outstripped our ability to create new names for new concepts. As a result, there are many networking terms which have confusing, double meanings. A simple glossary of terms does help to some degree; however, the terms are best explained in relation to each other. Many of the definitions, explanations and examples in this article will reference the sample network topology shown in the diagram below. [1, 2]

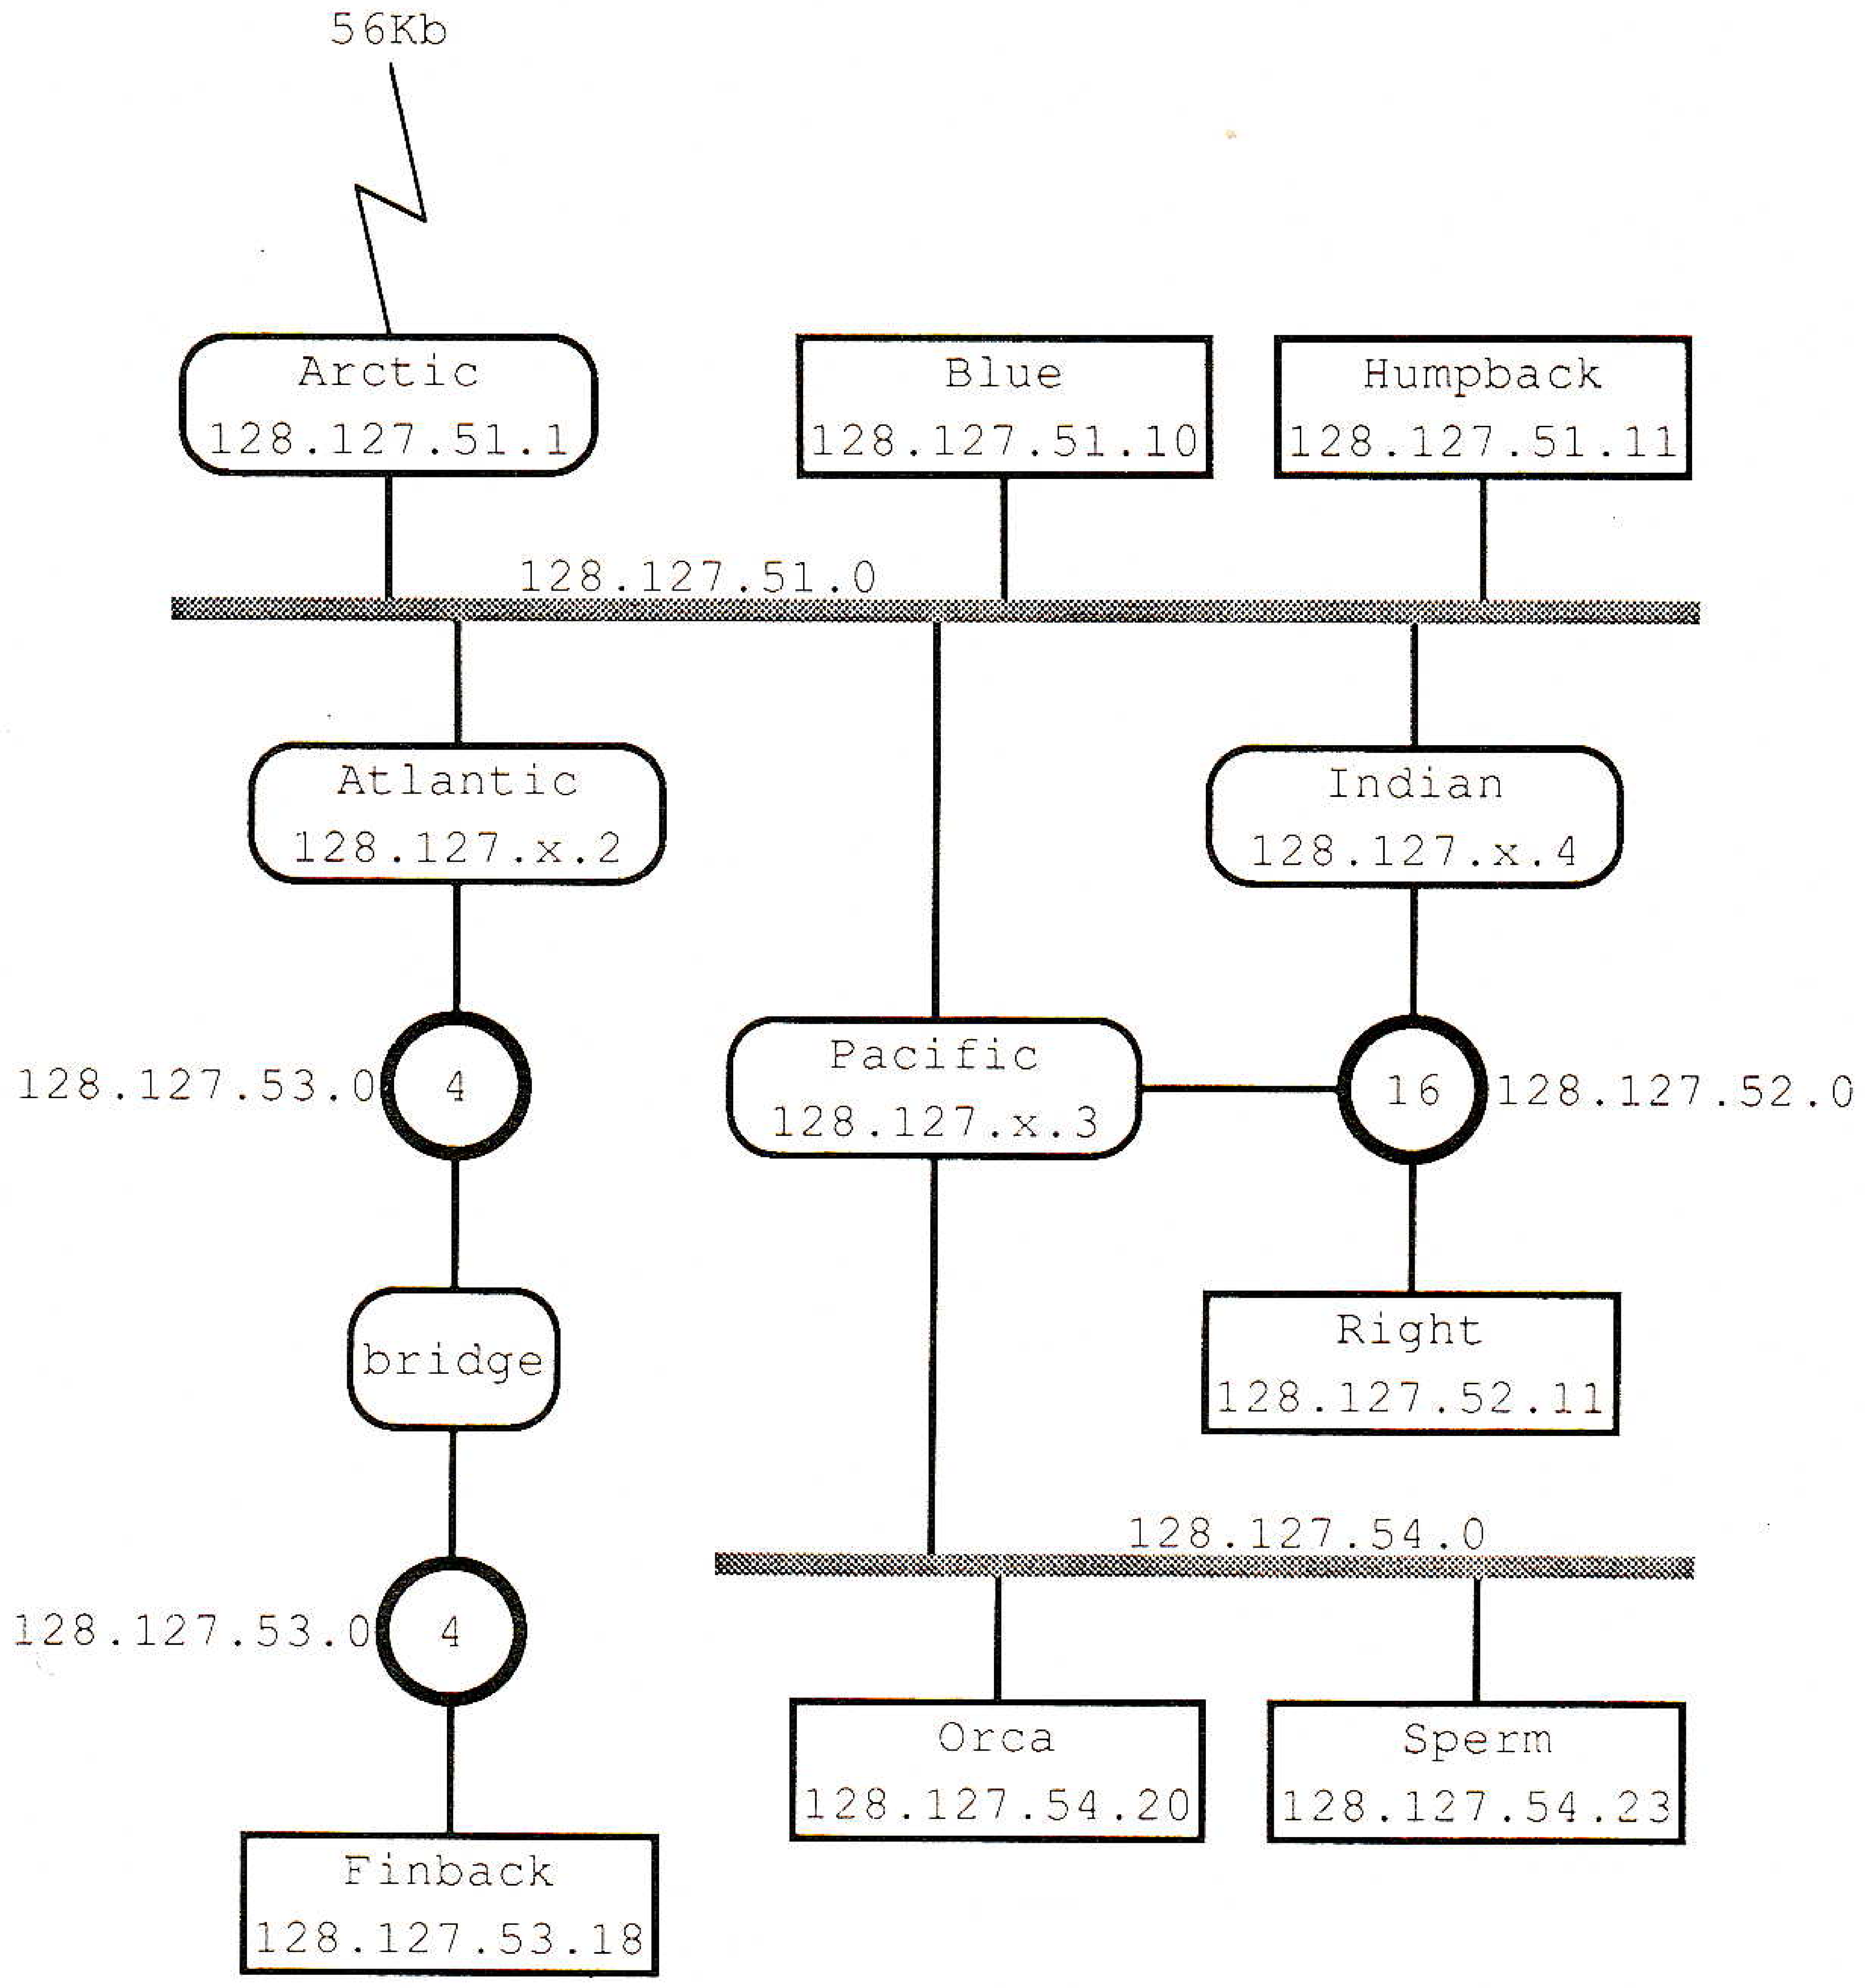

Nodes

A *node* is a device attached to a network. There are many types of nodes. For example, there are simple *repeaters* and *bridges* (which, being below users' attention thresholds, are sometimes not listed as "real nodes"); *routers,* which connect different networks and pass traffic between them in an "intelligent" manner; and *hosts,* which are user interface points to the network.

"Intermediate Systems" (IS) and "End Systems" (ES) are *Open Systems Interconnection* (OSI) terms for routers and hosts, respectively.

One also hears references to "gateways." Gateway is a much overused term in networking. It may refer to a router, a host which can perform routing functions, or an application layer router (e.g., a mail gateway).

Networks

Network is another much overused networking term. It is a generic reference to a context specific object. For example, when referring to a piece of an internet, it is common to say "network." However, when referring to a subnetwork of a network, it is also common to say "network." Although "network" is frequently used in a non-specific manner, it does have a very specific definition. A network may be defined as a collection of nodes and communication channels.

continued on next page

What's in a name? *(continued)*

For the purposes of this article, this is too general a definition. In dealing with the *Internet Protocol* (IP), the definition of a network is further restricted by specifying that the nodes share a common IP network address (more on addresses later).

Given the definition of a network, the terms *internetwork* and *subnetwork* can be defined. An internetwork, or internet, is a collection of networks (as defined under IP) joined together by communications channels. A subnetwork, or subnet, is a portion of a network which is independently addressable. A segment of a network need not be a subnet, and is not a subnet if nodes cannot uniquely identify it as being different than any other segment.

For example, consider the depicted network topology. Routers (named after oceans in this example) connect four different subnets (two Ethernet and two Token Ring) with six hosts (named after whales). Overall, it is a network, as defined above. There are four subnets (51 through 54). This network is also part of an internetwork, to which it is connected through the 56Kb serial line. Note that the two segments on the bridged Token Ring are not subnets since they cannot be independently addressed.

The Internet (proper noun) is an internet. It is the second largest network in the world (second only to the telephone system). It is composed of thousands of networks, with tens of thousands of nodes, and hundreds of thousands of users. It is used by governments, research institutions, educational institutions, and network vendors to send electronic mail and share information on virtually every subject known. [3]

Identifiers

Radia Perlman was kind enough to allow the use of her formal definitions for names, addresses and routes in this section.

If host X wants to communicate with host Y—

"Y" is a *name* if:
- "Y" continues to work, even if host Y moves; and
- "Y" works for any host X, regardless of host X's location.

"Y" is an *address* if:
- "Y" changes if host Y moves; and
- "Y" works for any host X, regardless of host X's location.

"Y" is a *route* if:
- "Y" changes if host Y moves; and
- "Y" is different for X's in different locations.

Using the oceanographic network, it is possible to codify the terms with concrete examples.

"Arctic," "Blue" and "Humpback" are examples of names. No matter to which subnet they are attached, those identifiers stay with the nodes. Also, all other nodes can use "Arctic," "Blue" and "Humpback" no matter where they are in the network. "128.127.51.1," "128.127.51.10" and "128.127.51.11" are examples of addresses. If Blue were to move to subnet 54, its address would change to "128.127.54.10." However, all other nodes can reach 128.127.51.11, no matter where they are in the network. A route is more difficult to exemplify since routes are very internal to the Internet Protocol and not generally visible to users.

IP Addresses

Consider, however, that a route from Finback to Humpback (the path which packets take) is different than the route from Orca to Humpback. Also, those routes could change if any of the involved nodes moved to another subnet.

An IP address is a 32-bit identifier, assigned to every node on an internet, which is guaranteed to be unique on that internet. The addresses is displayed, as seen above, as a dotted decimal number. Each portion of the number corresponds to an 8-bit byte. IP addresses are broken into four classes as shown in the table below. The table also shows the differences between the classes.

Class	*High Bits*	*Number of Nets*	*Hosts per Net*
A	0	126	16,777,214
B	10	16,382	65,534
C	110	2,097,150	254
D	1110	Multicast addresses	

As can be determined from the table, a class A addresses has one byte of network address and three bytes of host address. A class B address has two bytes of each. And a class C address has three bytes of network address and one byte of host address. In all cases, network and host addresses cannot have all bits 0 (meaning "this net") or all bits 1 (meaning broadcast). [5]

Subnetting

Subnetting is a way to add another level of interpretation into the addresses. It divides the host portion into two parts: the subnet and the host. The oceanographic network is a class B network (because the high two bits of the address are "10"). It has been subnetted by taking the high byte of the original host portion and calling that the subnet portion, and leaving the low byte as the new host portion. It is not necessary to divide on an even boundary. In fact, it is not strictly necessary to use contiguous bits; however, for user comprehension, it is recommended. [4]

References

[1] Quarterman, J., "Network Nomenclature," *ConneXions,* Volume 4, No. 9, September 1990.

[2] Lynch, D. C., & Jacobsen, O. J., "The INTEROP Glossary of Networking Terms," RFC 1208, 1991. [Based on "The INTEROP Pocket Glossary of Networking Terms," 1990. New updated version to be released at INTEROP 91].

[3] Dern, D., "The ARPANET is Twenty," *ConneXions,* Volume 3, No. 10, October 1989.

[4] *ConneXions,* Volume 3, No. 1, January 1989, Special issue on Subnets.

[5] Deering, S., "IP Multicasting," *ConneXions,* Volume 5, No. 3, March 1991.

GARY SCOTT MALKIN received his B.A. in Computer Science from Boston University in 1983 and expects to complete his M.S. in 1992. After almost five years working for Spartacus (he was the model for Fibronics' two advertisements), and two years at Proteon, he currently works at FTP Software in Wakefield. He is the principle software engineer for FTP's new router. Gary is also an active member of the Internet Engineering Task Force and the co-chair of the NOCTools Working Group. He is co-author of the "FYI on FYI" RFC and the "Questions and Answers" FYI RFCs.

Profile: THEnet

by Tracy LaQuey Parker, University of Texas at Austin

Introduction

The *Texas Higher Education Network* (THEnet) was formed in 1986 through a combination of networking efforts at Texas A&M University, the University of Houston, the University of Texas Health Science Center at San Antonio, and the University of Texas System. Covering the state of Texas, with a link to the Instituto Tecnologico y de Estudios Superiores de Monterrey in Monterrey, Mexico, THEnet connects over 60 academic and research institutions. THEnet's stated goal is to provide and advance the electronic exchange of information in support of the teaching, research, development and related collaborative activities of the Texas higher education and research communities.

Administration and Membership

THEnet was initially managed by the computing services staff of Texas A&M University. With the creation of the *University of Texas System Network* (UTSN) in 1986, the UT System *Office of Telecommunication Services* (OTS) networking staff assumed network management duties. OTS provides both network information center (NIC) and network operations center (NOC) services to THEnet member institutions.

Membership in THEnet is divided into three categories: Class A, degree-granting institutions of higher education, and their associated research institutions; Class B, nonprofit research and governmental organizations not associated with a Class A member; and Class C, industrial research organizations sponsored by a Class A member. Currently, THEnet consists of 40 Class A members, 11 Class B members, and 10 Class C members.

Protocols, Topology and Hardware

THEnet is not a homogeneous network utilizing a single networking protocol. Rather it is a network of physical connections between and within organizations making various use of IP, DECnet, SNA, NJE, OSI and compressed digital video to provide researchers, faculty and students the networking "tools" that they need for their particular situations.

THEnet utilizes Cisco Systems AGS routers between its hub sites at Dallas, Houston, Austin, and San Antonio. Other sites are connected to the hubs with a wide variety of equipment. Most of the data links are implemented with 1.5 Mbps leased digital data circuits. THEnet introduced its first full T1 data circuits in summer 1989, when in cooperation with *Sesquinet,* an NSFNET midlevel network administered by Rice University in Houston, it established a T1 triangle interconnecting Cisco routers located at UT Dallas, UT Austin, and Rice University. There are other T1 connections: UT Austin has T1 connections to UTHSCSA and M.D. Anderson Cancer Center in Houston; UT Dallas to UT Arlington and UT Southwest Medical Center; UTHSCSA to UT San Antonio and UT Pan American at Edinburg; and UT Pan Am at Edinburg to UT Pan American at Brownsville.

External Connectivity

THEnet's Internet connectivity dates back as far as 1973, when UT Austin established a connection to the original *ARPANET* [1] (the Defense Advanced Research Projects Agency Network). In April 1989, THEnet was officially designated an NSF midlevel network, gaining access to the NSFNET backbone through the *Nodal Switching Subsystem* (NSS) located at Rice University.

THEnet also provides access to the *Energy Sciences Network* (ESNET) through two backbone sites located at UT Austin and the *Superconducting Super Collider* Laboratory (SSC) in Dallas. ESnet is a multiprotocol (IP, DECnet and X.25) national backbone network operated and managed by the Department of Energy National Energy Research Supercomputer Center at Lawrence Livermore National Laboratory (LLNL). It provides connectivity for all DoE *Office of Energy Research* (OER) programs, including the *High Energy Physics* (HEP) community previously served by HEPnet (a DECnet network).

THEnet is active in national Internet activities, including membership in FARNET, a national organization of midlevel and regional networks connected to the Internet, and the *Internet Engineering Task Force* (IETF).

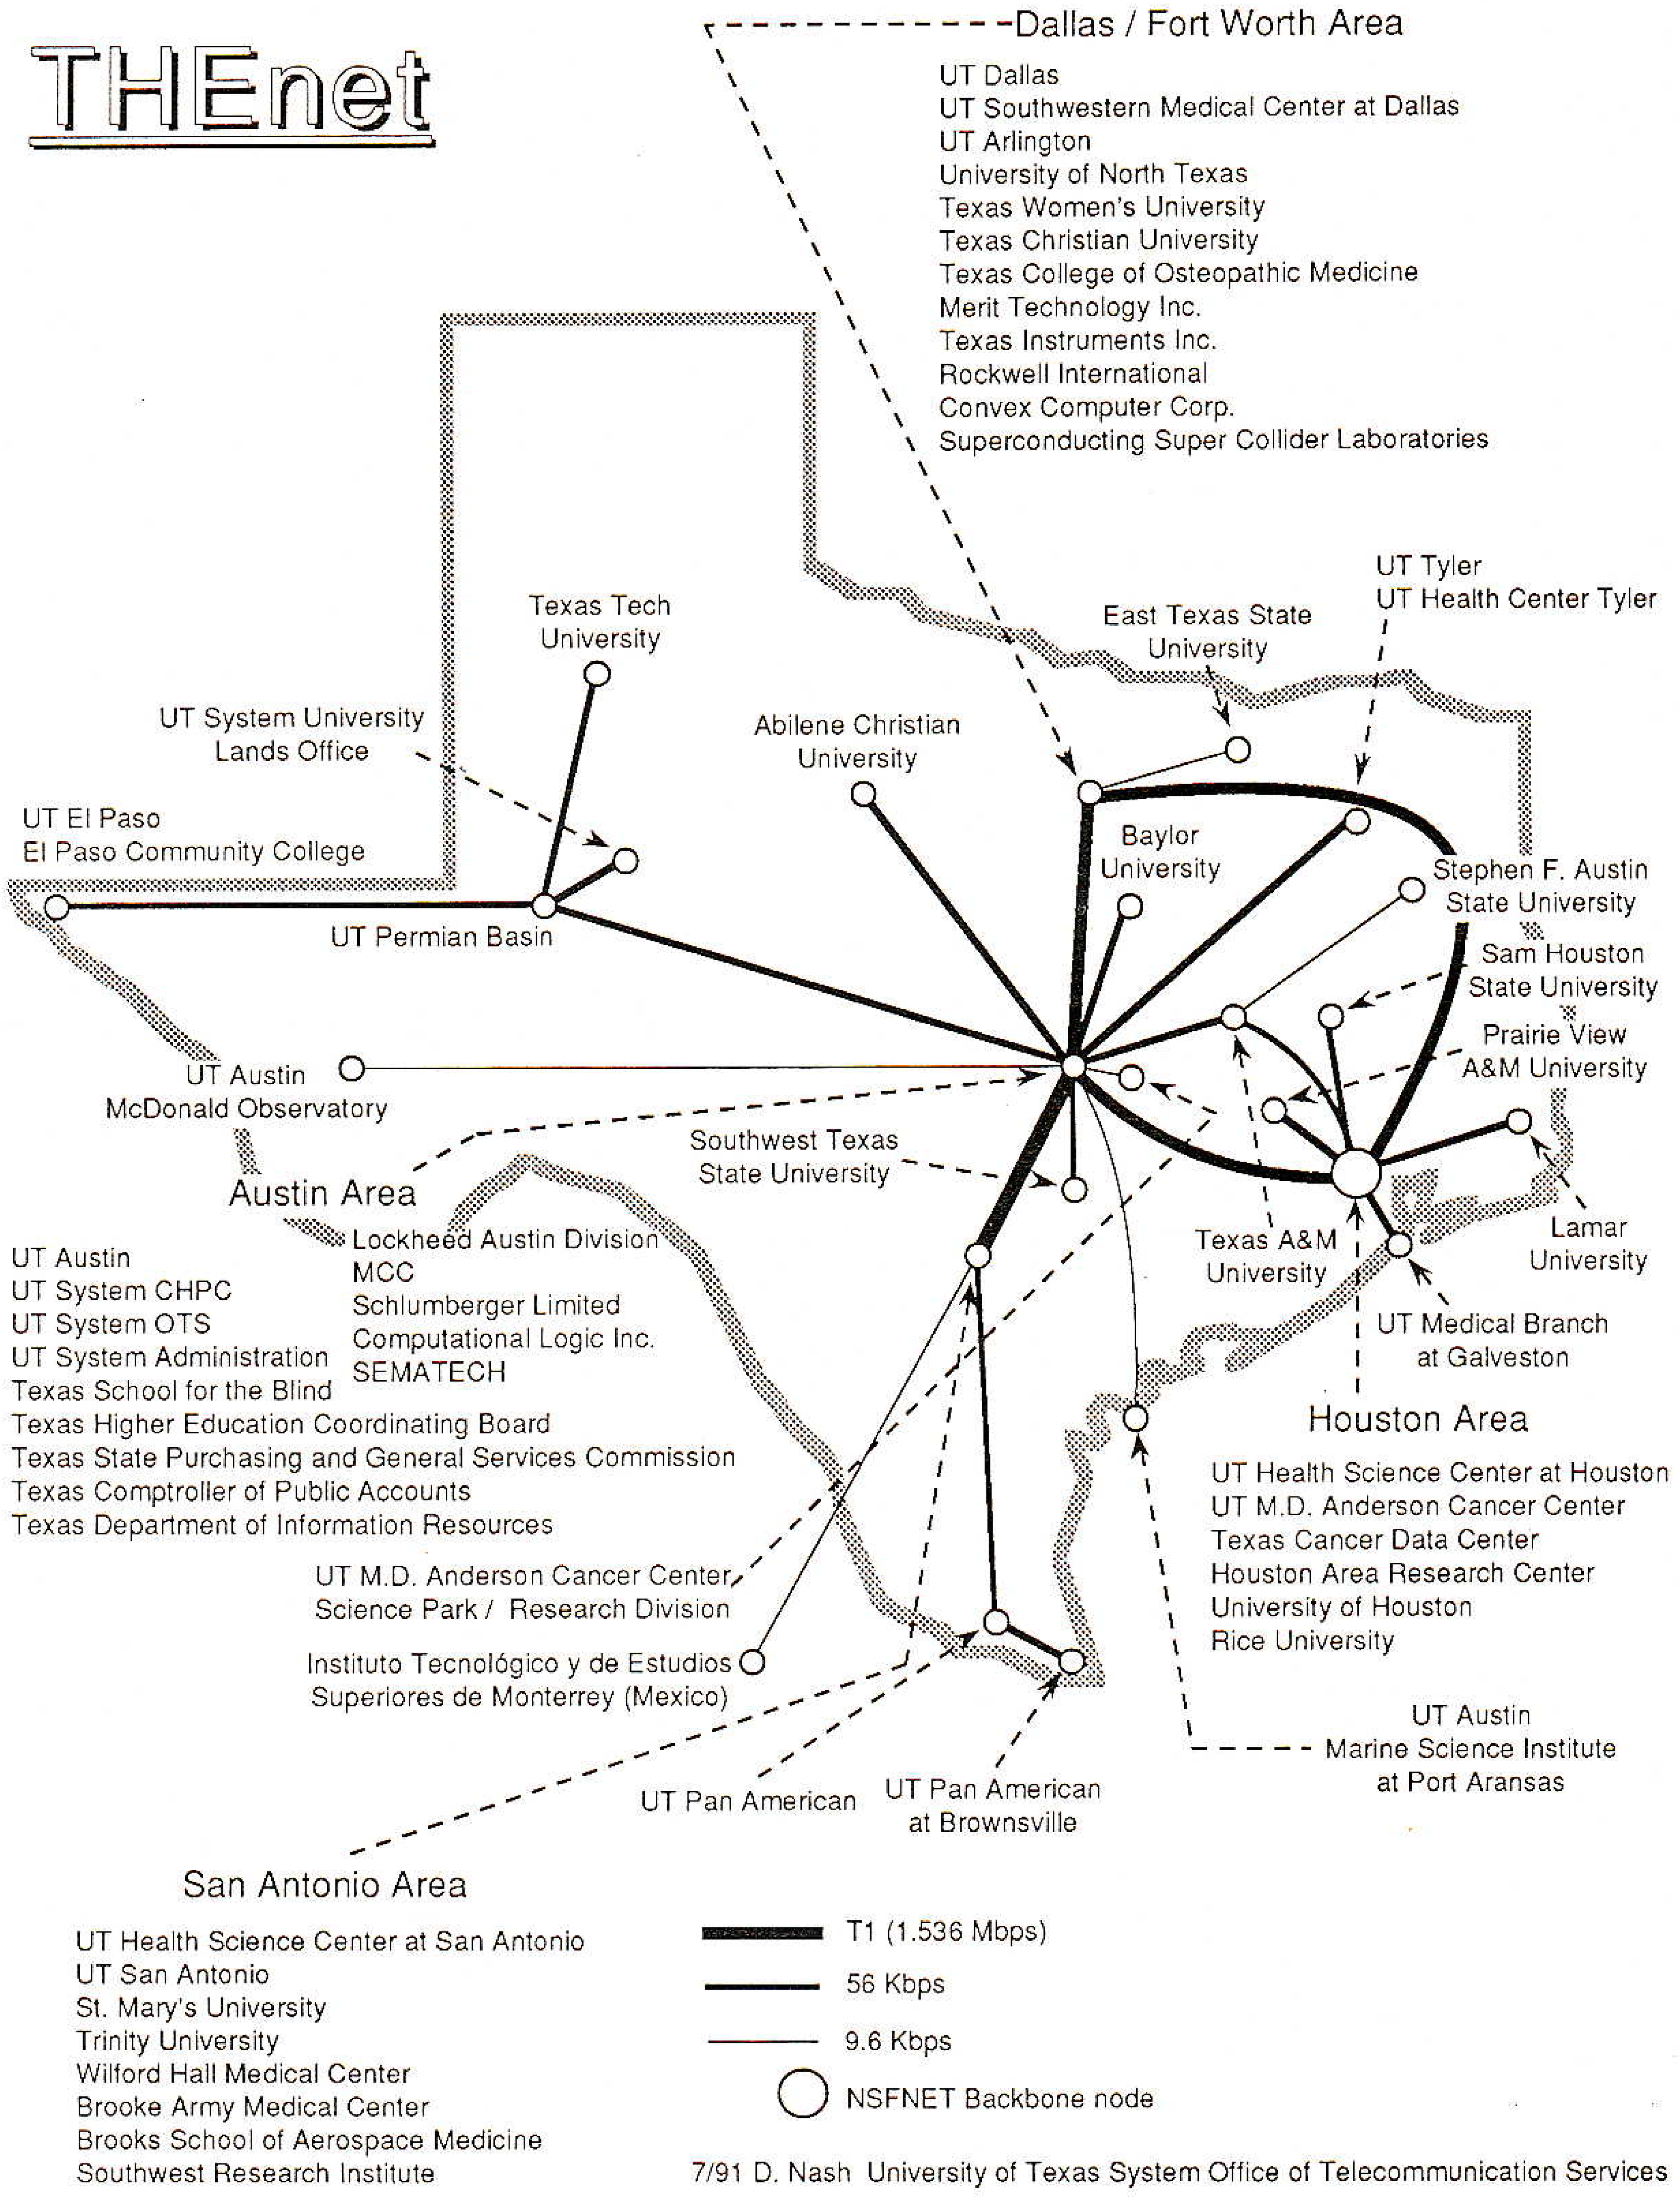

DECnet networks

THEnet provides *DECnet* (a proprietary protocol of Digital Equipment Corporation) [2] services for its members, and connects to two major nationwide DECnet networks. One connection is to the DECnet portion of ESnet and the *NASA Science Internet* (formerly the Space Physics Analysis Network, or SPAN) via DECnet routers located at UT Austin and the NASA Johnson Space Center. The DECnet address space of SPAN/NSI and ESNET conflicts with THEnet, so an address translation gateway (ATG) capability (which was developed at UT Austin) of Cisco Systems multiprotocol routers is used to selectively map nodes from one network into the address space of the other and vice versa.

continued on next page

Profile: THEnet *(continued)*

For nodes which are not mapped through the ATG, a technique known as *poor-man's-routing* (PMR) can be used to specify a path through a gateway system to reach a node on the remote network.

BITNET

A member of The *Corporation for Research and Educational Networking* (CREN) [3], THEnet offers BITNET services to many of its members and also participates in statewide BITNET network management. External BITNET connectivity is provided by Rice University via their participation in the BITNET II project (NJE over the TCP/IP based internet). Currently, the University of Houston (UH), UT Austin, UT Dallas, UT Arlington, and UT Health Science Center at San Antonio are all interconnected in a mesh topology with Rice University using TCP NJE. The other BITNET members in THEnet either partially participate in this mesh or are simply connected to one member of the mesh, using either NJE over TCP/IP or NJE over DECnet.

Public Data Networks

Gateway connectivity to the Telenet X.25 public data network is provided by OTS to THEnet members on a cost recovery basis. Currently, inbound X.29 terminal communication and VAX PSI mail gateway services are supported.

Network Services

Network information, consulting and operating services are provided through the THEnet NIC, which is run by OTS. Informative documents and host and contact lists are available on the THEnet NIC host, `nic.the.net`, via anonymous FTP, and `THENIC` (DECnet) via default DECnet file access (see the file `THENET.INDEX` for a list of available documents). Other useful networking information is also available for anonymous FTP on the host `ftp.utexas.edu` in the `pub` directory.

The *User's Directory of Computer Networks,* a directory of hosts, domains and contacts on major worldwide academic and research networks such as the Internet, BITNET, and the DECnet Internet, has been an annual publication of OTS since 1987. The 1990 edition was recently published by Digital Press.

In addition to NIC services, OTS hosts an annual telecommunication and networking conference. This year's conference will be held on August 20 at the Balcones Research Center in Austin. Speakers from the Computer Emergency Response Team (CERT), the Coalition for Networked Information (CNI), the National Agricultural Library, the Texas Education Agency (TEA), Advanced Network & Services, Inc. (ANS), Texas A&M University, VideoTelecom Corp., EDS, Thinking Machines Corporation, and The State of Texas House of Representatives, will make presentations about telecommunication and networking issues in higher education. For more information about future conferences, contact OTS (see address below), or send e-mail to `conference@utexas.edu`.

OTS also organizes tutorials on network configuration and management for THEnet members and THEnet Manager's meetings, during which technical and policy issues are discussed.

TENET

THEnet was recently chosen by TEA as the Internet services provider for its statewide K–12 communications network. This network, known as *TENET,* will provide Internet connectivity for at least 6,475 public schools by Fall 1991, and is expected to be fully operational during the 1991–92 school year.

TENET Topology

The initial installation will consist of high speed (V.32bis) modem pools established at THEnet points of presence in 15 cities. A toll-free 800 number will be available for schools not served by a dial pool in their local calling area. Users will be connected to SLIP capable terminal servers, which will in turn provide connectivity to the Internet and TENET's eight message processing and storage (MPS) computers located at University of Texas System components around the state. These UNIX-based computers will provide mail, conferencing, application, and disk storage services.

TENET Services

Educators will use Internet mail and USENET news for many projects, including curriculum-oriented conferencing and teacher development. Specialized information sources will be available either as locally provided databases and programs, or as remote services by third parties via the Internet.

There are many services that need development, such as better end-user applications for the PC and Macintosh, and cost effective ways to provide large numbers of LAN/WAN connections. THEnet will be actively developing software and services in these areas. Thinking Machine Corporation's *Wide Area Information Server* (WAIS) project is of particular interest and is currently being evaluated for use on TENET. [4]

TENET Information and Contacts

Further information about TENET is available via anonymous FTP on the host `ftp.utexas.edu` in the directory `pub/TEA`, or by contacting the following people:

Connie Stout, Texas Education Agency
`C.Stout@tenet.edu` • 512-463-9091

Jeff Hayward, OTS
`J.Hayward@utexas.edu` • 512-471-2444

Further information about THEnet

Queries about THEnet membership or additional information should be directed to:

Texas Higher Education Network Information Center
c/o University of Texas System
Office of Telecommunication Services
Commons Building, Room 1.156A
Balcones Research Center
10100 Burnet Road
Austin, TX 78758-4497

Voice:	512-471-2444 • FAX: 512-471-2449
Internet:	`info@nic.the.net`
THEnet (DECnet):	`THENIC::INFO`
BITNET:	`INFO@THENIC`
SPAN:	`UTSPAN::THENIC::INFO`

For further reading

[1] Dern, D., "The ARPANET is Twenty," *ConneXions*, Volume 3, No. 10, October 1989.

[2] Malamud, C., "DECnet/OSI Phase V: Real OSI or Only Selected Interfaces?," *ConneXions*, Volume 4, No. 10, October 1990.

[3] Laubach, M., "Profile: CREN—The Corporation for Research and Educational Networking," *ConneXions*, Volume 4, No. 5, May 1990.

[4] Kahle, B., "An Information System for Corporate Users: Wide Area Information Servers," to appear in *ConneXions*, 1991.

TRACY LaQUEY PARKER received a bachelor's degree in Computer Science from the University of Texas at Austin in 1986 and currently works at the UT Austin Computation Center Networking Services Group as a Network Information Specialist. She is a member of the IETF and also represents THEnet in FARNET. She is the editor of the *User's Directory of Computer Networks*, published by Digital Press.

Research and Academic Networking in South Africa

by Mike Lawrie, Rhodes University

South Africa

The country is located at the Southernmost end of the African continent. The time zone is two hours ahead of UT. The country lies between the tropic of cancer and 34 degrees South—i.e., no further from the equator than South Carolina or Baghdad. Population is approximately 35 million, 85% of which are black. The country is very roughly 1000 miles by 1000 miles square.

Mineral wealth (gold, diamonds, coal, iron ore) has allowed a good infrastructure of roads and telephone services to be established between the coastal cities and the industrial regions in the North of the country. The country can feed itself.

South Africa is currently emerging from the apartheid era, and needs to use technology to uplift those who have been disadvantaged by this inhuman policy. Networking is very important, because without e-mail facilities the chances of attracting good academic staff from the First World are severely reduced.

There are two official languages, English and Afrikaans. There are a variety of African languages, of the Nguni stem, of which Zulu and Xhosa (say "korsa" and you are only slightly wrong) are widely spoken. In most parts of the country you will be able to converse in English.

History

The first computers were installed in Universities in South Africa in the mid 1960s. An overly optimistic idea in about 1967 to connect these computers into a network died a quick death—data communication alone was a stumbling block, quite apart from the attitudes in the minds of the role-players.

The *Council for Scientific and Industrial Research* (CSIR) was an early user of computers for research, and had a network in place by the 1970s. This was based on Case multiplexors, and ran multi-protocols (async, IBM SNA). This allowed the CSIR research institutes to access computers located in Pretoria. Several Universities had access, but only the University of Natal and Rhodes University had their computers hooked onto this network, and even then this was not until the 1980s. The network was primitive, and did not encourage file transfer between host computers. E-Mail was barely possible—users had to log into one central computer to use e-mail, and this did not encourage usage.

Several (but not all) universities in the early 1980s had a local network to serve their campus, but even now (1991) backbone LANs are not installed at all sites.

University links

In the early 1980s, three universities about 80 miles apart (Witwatersrand, Pretoria and Potchefstroom) hooked their IBM computers into a network, using dedicated lines at 9600 bps. This gave the standard facilities of IBM's VM system. By some accounts, not a great deal of use was made of this network, but it was a start. One of the CSIR computers was hooked in, as well as some computers of the Government network known as *Govnet*—this included the *Sabinet* library system. All hosts were IBM, using SNA.

In 1985, Rhodes University coupled its CYBER to its VAX using JNET, and then linked via a 600-mile dial-up 2400 bps connection to this network. This was solely for the purpose of e-mail (using a home-brewed mailer that could be hooked into CDC's NJEF package). The style of working was similar (we believe) to BITNET.

Feelers had been put out to the international networks, to see whether links could be established. Regrettably, efforts were not co-ordinated, and the country was by now (1986) in the grip of an effective international sanctions campaign. The attitude of some of the South African networkers was that it was a waste of time to try to connect. This was borne out by the experience of that very liberal institution, the University of Cape Town, which had a link for about three weeks to Uunet before the connection was summarily cut.

Fidonet

Fidonet has been operating in South Africa since 1986. This was initially a network of 5 PCs, using dialup circuits. International links were established via Europe early in 1987.

Following on from some helpful suggestions given to Rhodes University by `hostmaster@sri-nic`, Fidonet appeared to be the best option to follow in order to get communication to the Internet.

In 1988, Rhodes University set up a PC to link into the international Fidonet network, and wrote an interface so that the home-brewed mailer on the Cyber could exchange mail with Fidonet. This allowed Cyber users to access the Fidonet network, which in turn could access the world's e-mail networks. A minor extension allowed all of the computers on the South African network to get at international mail, and once the political issues had been addressed (mid 1989), researchers in South Africa at five institutions (Rhodes, Witwatersrand, Potchefstroom, CSIR and Pretoria) could get international e-mail. By an enormous margin, the first two of these were the biggest users.

By deft use of *Kermit,* computers at UCT and the University of Natal joined in by connecting to Rhodes, but reliability was poor. These links were soon converted to TCP/IP, and this was the start of the internet on Uninet.

Yet another Uninet is born

Late in 1987, the *Foundation for Research Development* (then a part of the CSIR) had taken an interest in networking developments. The universities themselves were unable to get the basic infrastructure in place. The FRD set up a project called *Uninet,* to establish a research network in South Africa. A study was commissioned (it missed out recommending TCP/IP!), and in due course a heavy investment in trunk circuits and multiplexors was made. This catered for several protocols, and had a degree of flexibility.

Protocols catered for were RSCS, DECnet, Case proprietary and fortunately async operation (used by PCRoute). The policy now is to "get to the standards of the Internet, and to follow those standards." The stage was now set to let the Universities grow the network on top of this infrastructure.

Control structures and interested parties

The biggest single stumbling block to any networking in South Africa is the "third party traffic" regulation of the South African Posts and Telecommunications body. Not even senior staff have a consistent interpretation of this rule, but the effect on users is that they are very wary indeed about allowing any form of data communication on behalf of other organisations, in particular forwarding e-mail. Universities may now forward each other's mail, having got special clearance from the SAPT to do so, but it is going to be very difficult to have communication linked to commercial organisations (e.g., vendors who could provide support to the Universities).

continued on next page

Networking in South Africa *(continued)*

There are about twenty universities and five research institutes that are associated with the Uninet network, and the number is growing. South African universities are autonomous institutions, and certain of them have fought hard to maintain that autonomy during difficult political circumstances. International contact was, at times, very difficult indeed, and at the same time there was tremendous pressure from the government to toe the line "or else." All South African universities are heavily dependent on annual grants from the government in order to operate.

The Universities have a voluntary association called the *Committee of University Principals* (CUP), that spawns further committees. One such committee, on computing, made valiant attempts to get networking going, but was hamstrung by funding and feuding—hence the efforts of the FRD were welcomed.

The Foundation for Research Development is the body that co-ordinates university research. It uses a peer-evaluation system to assess the abilities of research workers, and assists researchers with funding according to their rating. While originally part of the CSIR, the FRD is now an independent organisation. Having laid out the seed money, the FRD appears to be reluctant to release control of the network. This might or might not be a good idea.

The *Uninet Control Board* is an advisory panel that meets at most twice a year. Three members are elected by the participants of Uninet, three are appointed by the FRD, and a chairman is appointed by the FRD from among the six members; the Uninet Manager is ex-officio a member. Day to day control of Uninet is done by the Manager, who has a technical assistant. The Manager is a part-time employee of the FRD. The FRD underwrites the operation and, in principle, accepts responsibility for the managerial and administrative costs, while the participants are expected to bear the operational costs.

Administration of the Diginet (see below) multiplexors is contracted to the CSIR. Operation of the international gateway is controlled by Rhodes University.

All data communication across property boundaries is the monopoly of the *South African Post and Telecommunications* (SAPT). Satellite technology is not permitted, and there are very tough rules about the use of dedicated circuits—in particular relating to the forwarding of traffic on behalf of other organisations. Universities are not exempt from this "third party rule," and this is having an inhibiting effect on expanding the network to allow commercial organisations, such as computer vendors, to join (i.e., a `.com` domain) and to share the costs.

Internal links

There are about 4500 miles of digital trunks connecting major centres. These operate at 64 Kbps on the SAPT *Diginet* network—a fibre-optic and microwave system. Multiplexors allow static partitioning of the trunks into independent channels. Physical routes radiate out from the Pretoria/Johannesburg area to the coast, and there is very little in the way of redundancy. The logical routes result sometimes in a message passing through one site up to six times on different protocols. As the TCP/IP usage grows, this will be rationalised, and the total trunk distance could be halved.

Due to the variety of machines on Uninet, there are currently a variety of protocols on the trunk circuits—e.g., TCP/IP, RSCS, DECnet and even Kermit.

Uucp is used on a few dial-up links. There is no doubt whatsoever that the TCP/IP usage is experiencing a strong growth. The other protocols appear to be entrenched, and this is costing Uninet twice as much for the trunks compared to a pure TCP/IP environment. Most of the protocols run at 9600 bps, but most of the TCP/IP routers are operating at 19,200 bps.

TCP/IP routing is done with Vance Morrison's *PCRoute*. This is for a variety of reasons, most important of which were availability and cost. Approaches made to vendors of routers met with a polite "Sorry, you're from South Africa, no deal." (sigh). In due course purpose-built routers will be used, and these will run on direct circuits of at least 64 kbps. Currently there are six routes in operation, with two more about to start up.

Uucp is used to link into a loose association of netters called *SANET*. This association has led to an interesting situation, perhaps unique in the world, viz:-

All along the way, great care has been taken to avoid compromising the e-mail gateway due to certain provisions of the *Comprehensive Anti-Apartheid Act* (CAAA). This act prohibits persons in the USA from having dealings with certain organisations in South Africa—e.g., military, police and "apartheid-enforcing entities." It does *not* prohibit contact with South Africa, but many places do not wish to be seen to be associated with this country, due to commercial pressures in the USA. But there are individuals and organisations in the USA that see beyond the immediate problems, and believe that communication is more likely to assist a transition than to hamper it. Also, some organisations within South Africa (e.g., some universities) will themselves not deal with organisations covered by the CAAA. Hence, at all times there has been a political sword hanging over the network, which has added a dimension not normally associated with networking.

International links

The original international Fidonet link is still in place. This has grown to become the zonegate between Fidonet zones 1 and 5 (USA and Africa). It uses a 386 system, called *Settler City*. The protocols used on Fidonet are extremely efficient, and after the initial bedding in of the procedures, the system has served Uninet very well indeed. It may well be nearing the end of its useful life, however.

The kludge used to route e-mail through Fidonet relied on having a flat BITNET style address space on Uninet. With the Fidonet address of Settler City being 5:4/494, the Fidonet/Uninet gateway at Rhodes University translated between an early-style Uninet address of:

`user@host`

and a domain address of: `user.host@f4.n494.z5.fidonet.org`.

Outgoing mail gets forwarded to the zone 1 zonegate, from which it gets into *uucp* via a Ufgate (*uucp* to fido gate) system, and hence to the world's networks. There are Fidonet links to Botswana, Zimbabwe, Kenya and Ethiopia, with the latter not being particularly reliable. Volumes are not high.

Uucp

In June 1990, a *uucp* link was operational. This uses a V.32 modem, and connects via `m2xenix` in Portland, Oregon. Throughput is very disappointing, being about a third of that of Fidonet, but it does have the advantage of better address space, and uses a more rugged network. In addition, access to newsfeeds is simplified. Hardware is a 386 with 300 Mb disk, software is SCO Xenix with TCP/IP.

continued on next page

Networking in South Africa *(continued)*

This second gateway at Rhodes presented a problem, as it was now necessary to maintain compatibility with the Fidonet system, yet at the same time to take advantage of the extra features. There was a conflict with the flat namespace of Uninet, and the desire to move on to the DNS that was now possible with the fledgling internet. The DNS system won.

During a visit to INTEROP 90 in October, contacts were made and procedures sorted out for the registration of the `.ZA` domain. This was functioning by December 1990.

International e-mail traffic to South Africa runs at about 10 Mbyte a day, and is extremely reliable.

There are *uucp* links to Namibia, essentially only to one individual. The `.NA` domain has recently been registered, and mail will flow via the Uninet gateway at Rhodes University. A link to Zimbabwe is on the drawing board. Enquiries have been received for links into other parts of Africa, and Uninet is keen to assist where possible.

The future

IP connectivity is not very far away. This will be of great benefit, as many South African researchers have good relationships with their overseas counterparts, and being able to access remote computers will be a breakthrough. A 9600 bps link is about all that can be afforded, but it is likely to be overloaded.

Connections within the country will move away from the vendor protocols such as DECnet and SNA. TCP/IP will take over, and direct trunk circuits will be installed, with better resilience against failure of routes. These circuits will be at least 64 kbps. Purpose-built routers will take over from PCRoute.

There will be better and more links into other parts of Africa. Already the political logjam is clearing, and contacts with Zimbabwe and Namibia will see rugged links appear in the near future. The biggest problem in these countries is to get the ball rolling, and expertise is needed more than anything else.

There is no reason why anyone in a low-technology country of Africa cannot get e-mail via Uninet. All that is needed is a telephone, a modem and a PC with Fidonet software (which is public domain). The Fidonet protocols are extremely rugged, and put *uucp* to shame with regard to performance and reliability. The Uninet project is keen to co-operate with potential e-mail users in Africa, and the skills and experience that have been acquired at Rhodes University are for sharing.

Stop Press!

The South African research network will be IP-addressable as soon as the dedicated line to the USA is operational. Expected date: 8/1991.

Acknowledgements

Comments and ideas from Randy Bush (Portland, OR), Vic Shaw (FRD, Pretoria) and Henk Wolsink (Fidonet, Port Elizabeth) have been incorporated in this article.

MIKE LAWRIE `<ccml@hippo.ru.ac.za>` qualified as a computer technician with ICL (SA), went on to get a BSc (Hons) degree at Rhodes University, and subsequently a Masters degree at UMIST in Manchester, UK. He is currently Director, Computing Services at Rhodes University, and has been responsible for computing there since 1971. He served on the original university committee investigating networks, and serves on the Uninet Control Board.

[Ed.: Mike Lawrie will speak at the International BOF at INTEROP 91. The BOF is Wednesday, Oct. 9 from 6pm to 8pm].

Call for Participation

The *Symposium on Experiences with Distributed and Multiprocessor Systems III* (SEDMS III) will be held March 26–27, 1992 in Newport Beach, California. The symposium is sponsored by The USENIX Association in association with The Software Engineering Research Center (SERC) and in cooperation with ACM SIGARCH, SIGCOMM, SIGOPS and SIGSOFT IEEE-CS Technical Committees on Distributed Processing, Operating Systems, Software Engineering, and Design Automation.

Goal

The goal of this symposium is to bring together individuals who have built, are building, or will soon build distributed and multiprocessor systems. SEDMS III will provide a forum for individuals to exchange information on their experiences, both good and bad, including experiences with coding aids, languages, debugging and testing technology, reuse of existing software, and performance analysis. The presentations should emphasize the lessons learned from use of such systems and tools.

Format

Extra-long breaks between sessions and work-in-progress presentations will be provided to facilitate a workshop-like atmosphere during parts of the symposium. We will also have discussion panels on submitted themes.

Submissions

Six copies of each submission or panel proposal should be sent to the program committee chair (address below) to arrive no later than 1 November 1991. Submissions of full papers are invited on any topics related to the theme of the symposium. The committee will give preferential consideration to submissions describing experiences with actual systems-papers describing purely theoretical work will not be accepted. Panel proposals should include a description of the relevance to the goals of the SEDMS, and the qualifications of the participants suggested.

Important dates

Submissions due:	1 November 1991
Notifications mailed:	20 December 1991
Camera ready copy due:	24 January 1992

More information

For Further Information, or to have your name added to the mailing list for registration materials, send your name and surface mail address to either the general chair or program chair:

George Leach, General Chair
AT&T Paradyne
MS LG-129
PO Box 2826
Largo, FL 34649-2826
813-530-2376
`reggie@pdn.paradyne.com`

Gene Spafford, Program Chair
Software Engineering Research Center
Dept. of Computer Sciences
Purdue University
West Lafayette, IN 47907-1398
317-494-7825
`spaf@cs.purdue.edu`

Call for Submissions—*IEEE Network*

Background

IEEE Network Magazine is the IEEE's flagship publication in the area of computer communication. To better serve the interests of the of the computer communication community, *IEEE Network* is going through an editorial restructuring under a new Editor-in-Chief, starting with the January 1992 issue.

One of the new editorial goals is to organize the issues of your magazine around topics perceived to be important by the computer communications community at large. This involves organizing the issues around articles received through open submission on topics of interest to the community. Good quality submission from you will be the key to the Magazine's success.

IEEE Network publishes technically refereed articles on all topics related to the exchange of information—of any type, including data, text, images, and voice—among computers or among computers and individuals using computer based communication technology. To help ensure that good submissions are published in a timely fashion, the editorial board's goal is to return reviews to authors within ten weeks of submission and to publish articles within six months of receiving final copy.

The editorial goal of the magazine is to provide its readers with highly readable, carefully edited, useful, and timely information. This includes articles discussing new results and technologies as well as articles on useful practices and techniques, retrospective articles, and commentaries or reasoned polemics. As an illustration of the general flavor desired, a submission that sketches the underlying mathematics of a new queueing theory result and then illustrates its usefulness by presenting an alternate solution to a network planning problem would be highly preferable to one that dwells on the mathematical development of the result itself.

Topics

Topics of interest include, but are by no means limited to, the following:

- Progress in the development and implementation of gigabit and terabit data networks and switching systems;
- Naming and directory services, especially for systems designed to scale to billions of end systems;
- Traffic analysis and modeling techniques;
- Network operations;
- Novel distributed applications;
- Thoughts about the future of communication networking (integrated services? to what extent? why?);
- Synchronization (e.g., voice and video channels);
- Network security (both practices and mechanisms);
- Technology deployment and enhancement
- Solutions to deployment problems (such as getting FDDI interfaces to got at 100 Mbs; implementing 802.6 interfaces; etc.)

Submissions

Prospective authors encouraged to show their interest in and enthusiasm for computer communications by submitting articles on topics of their choice to the new editorial board at the following address:

IEEE Network Magazine
Attn.: Editor-in-Chief (submissions)
IEEE
345 East 47th St
New York City, New York 10017-2394
USA

Please do not send submissions to any other address, as only submissions sent to the IEEE in New York will be considered.

A method for submitting articles electronically is in preparation. Contact us at the address above if you are interested in submitting electronically.

Craig Partridge, Editor-in-Chief (effective January 1992)
John Daigle and John D. Spragins,
Senior Editors IEEE Network Magazine

SRI International announces TCP/IP CD ROM

The first TCP/IP networking CD is now available through SRI International's Network Information Systems Center. Designed for users and implementors of TCP/IP networks, the CD-ROM contains an easy-to-use search program for protocol standards, informational documents, and historical networking archives as well as networking source code.

Content

Specific information provided by the CD includes the Internet *Request For Comments* (RFCs) series of documents, *Internet Engineering Notes* (IENs), and the document collection from the DDN Network Information Center. As a free bonus, it includes source code related to TCP/IP networking for those users without easy access to the Internet archives.

Using indexes already stored on the CD, a simple search program quickly finds references to keywords. Indexes have been generated for the RFC and IEN collections as well as the Internet *TCP-IP* and *Namedroppers* (domain naming) mailing list archives.

Format

The TCP/IP CD is an ISO-9660 (High Sierra) format compact disk. Files are formatted for UNIX and MS-DOS systems and may be readable on other systems supporting ISO-9660 file systems. It is priced at $395, which includes one free update disk.

More information

To purchase the TCP/IP CD or for more information, contact:

SRI International
Network Information Systems Center – EJ291
333 Ravenswood Avenue,
Menlo Park, CA 94025
Phone: 415-859-NETS (That's 415-859-6387)
Fax: 415-859-6028
E-mail: `TCP-IP-CD@NISC.SRI.COM`.

SRI International provides research, development, and consulting services to business and government worldwide.

[Ed.: We hope to have a "CD Review" of this CD in a future issue.]

Gore introduces Fiber Optic Network Bill

Incentives to develop the fiber optic communications network critical to U.S. economic strength in the Information Age, are key to legislation introduced in June by Senators Al Gore, D-TN, and Conrad Burns, R-MT, that sets strict guidelines to allow telephone companies to provide cable television services.

"The strength of our economy increasingly depends on our ability to move information and our ability to move information depends on the quality of our communications networks," said Gore. "As a nation, we must act now to advance the construction of these networks. Our foreign competitors are moving forward. We cannot be left behind."

The *Communications Competitiveness and Infrastructure Modernization Act of 1991,* introduced by Burns and Gore, provides incentives for modernizing U.S. telecommunications systems and, over time, removes restrictions that now prevent telephone companies from providing cable television services. The bill sets 2015 as a goal for the development and use of these advanced fiber optic networks.

Fiber to the home

"By 2015, the Japanese plan to have connected every business, home, and institution served by these fiber optic networks. If the United States fails to update current policy, we will be decades behind and the price we will pay—in jobs, in innovation, in medical and other research—will be enormous," said Gore.

Gore is the author of two related bills, one to increase competition in the cable television industry and a second to develop a high-speed fiber optic network to link the country's most powerful computers.

Real competition

"While cable television companies enjoy all the benefits of a monopoly, cable customers suffer all the ills: skyrocketing rates, poor service, no choices. In rural areas where cable companies can't find big profits, sometimes there's no service at all. Real competition could reduce cable television rates by one-half and save consumers $6 billion a year," Gore said, referring to a study by the Consumer Federation of America.

Video dial tone

The bill would initially allow telephone companies to transport video programming offered by other companies, what is referred to as a "video dial tone" service. Then, with state regulatory commission and FCC approval, the telephone companies could offer their own programming on 25 percent of the available channels, leaving 75 percent available for others who provide programming.

Highlights

The bill was introduced June 5, 1991 in the U.S. House of Representatives by Representatives Rick Boucher, D-VA, and Michael Oxley, R-OH. The legislation (S. 1200):

- Sets a new national goal of 2015 for the U.S. to establish an advanced, interactive, interoperable fiber optic system available to all homes, businesses, educational institutions, health care organizations and other users;
- Requires telephone companies to submit to their state regulatory commissions a construction and deployment plan that would include a schedule for completion of the system by 2015, a description of the technology to be used, estimates of construction cost, and the plan for retirement of plant displaced by the new system;
- The telephone companies must each submit a plan that provides for early use of these networks by educational institutions, health care facilities, and small businesses and, at a reasonable pace, must also include less densely populated areas and economically disadvantaged areas;

- Each state commission will have one year to act before the FCC will review the plan to certify compliance with national goals and the objectives of the Burns-Gore bill;
- Telephone companies would initially be precluded from providing cable television services under strict guidelines until after the state and the FCC approve the telephone company plans and then, the telephone companies could provide cable television services under strict regulatory safeguards;
- The strict guidelines would prevent telephone companies from subsidizing cable television services with telephone rates and prohibit the buy-outs of existing cable systems. The bill also includes a "death penalty" provision that, if these guidelines are violated, would force the telephone company to divest itself of its cable business. And, the legislation prohibits cross-marketing of video and telephone services.

TCP/IP Traffic Study paper available

A new tech report (LBL-30840): *Measurements and Models of Wide Area TCP Conversations,* by Vern Paxson, LBL Computer Systems Engineering Group is available via anonymous FTP from host `ftp.ee.lbl.gov` (128.3.254.68), file `tcpmeasurements.ps.Z` (this is a 1MB compressed *PostScript* file).

Two months of data collection

This study is similar to that described in the recently announced paper by Danzig, et. al., of USC. However, a much more efficient means of data capture was employed (described in the paper) which allowed the collection and analysis of every conversation into and out of the LBL campus over a period of two months, rather the the two hour period of the USC study. This greater temporal scope allowed the investigation of long-term TCP/IP traffic variation (i.e., hourly, daily, weekly and monthly). It also allowed a detailed investigation of the geographic distribution of traffic. (This and other studies have shown that there is substantial short term correlation in conversation patterns and it is not possible to reliably assess geographic distribution from samples spanning less than a few days.)

Abstract

"This paper describes measurements of all of the wide area network TCP conversations between the *Lawrence Berkeley Laboratory* (LBL) and the rest of the world for the months of November, 1990, and March, 1991. Some 500,000 conversations were recorded, encompassing 11 different major protocols. We look at aggregate characteristics of these conversations, both overall and by TCP protocol (e.g., SMTP, FTP), computing the distributions of amount of data transferred, network bandwidth used, conversation lifetimes and conversation interarrival times. Temporal traffic variation is also investigated, showing the variation of number of active conversations and network bandwidth utilization over periods of 24 hours, 7 days and 30 days. Long term variation is also investigated by separately analyzing November and March data (which reveals a 10–20% increase in almost all aggregate traffic characteristics in just four months). We classify each conversation geographically and discover that the connectivity of the conversations was remarkably rich, including traffic to 48 of the 50 states in the U.S. and 23 foreign countries. Finally, we develop a number of models for describing conversations of the various protocols. From these models we can more readily assess how each protocol is used and how the use changes as network utilization grows."

Solutions Showcase Demonstrations for INTEROP 91 Fall

Witness exhibitors cooperatively demonstrating the interoperability of their products and technologies—the proof you need to make informed buying decisions. Our INTEROP 91 Fall Solutions Showcase Demonstrations feature these technology areas:

FDDI

For the third consecutive year, INTEROP 91 Fall takes the lead in showcasing *Fiber Distributed Data Interface* (FDDI) products for an increasing array of high-performance network computing applications. FDDI is a high-speed, 100Mbs, general purpose LAN standard, originally designed for optical fiber, but extensible to support alternative media. This year the focus will be on FDDI in real-world environments, with several special interest groups for emerging FDDI applications in client-server computing, campus backbone networks, and network management tasks. Over 30 demo participants are committed to this "watershed year" for FDDI.

Frame Relay

Frame Relay is a wideband data service that satisfies today's requirements for wide area networking. Supported by many hardware vendors and service providers, it has applications in LAN/WAN internetworking and access. Multiple private virtual circuits can share a single high-speed access line allowing full mesh connectivity to multiple locations. Network delay is substantially lower than with X.25 networks. The showcase will demonstrate switching equipment, access devices, and public carrier services working together to provide an intelligent data network. It emphasizes the interoperability of many products over wide area network facilities. LAN bridging and routing, as well as terminal to host applications will be demonstrated.

ISDN

On the verge of implementing National ISDN 1, it is timely and appropriate that the INTEROP 91 Fall *Integrated Services Digital Network* (ISDN) showcase will demonstrate multivendor, interoperable, multimedia communication services employing basic and primary rate interfaces (as defined by CCITT international and Bellcore national standards). The showcase will demonstrate the power of ISDN as a global, ubiquitous telecommunications service infrastructure—a revolutionary step beyond today's single-vendor, proprietary island deployments. True ISDN is about to arrive!

ONC/NFS

The trend toward distributed, heterogeneous computing is very clear. Users are increasingly demanding transparent information access—without regard for the specific hardware, operating system, or network architecture in use. Ultimately, network computing must bring all organizational computing resources to the user's fingertips. INTEROP 90 hosted one of the first public multivendor demonstrations of one such enabling technology—ONC/NFS. (*Open Network Computing/Network File System*). The ONC environment includes two major functional areas: *Distributed Resource Access,* including NFS which provides transparent access to remote files over a network, and *Distributed Application Development,* including *Remote Procedure Call* (RPC) which provides a high-level mechanism for executing operations on remote systems. Applications included in this year's demonstration will span all industry segments, including Financial/Accounting, Inventory Management, Office Automation, Medical Systems, and Network Management. In addition, a section of the demonstration will be dedicated to explaining ONC technology and how it is used by software developers to build distributed applications.

OSI

Continuing the tradition of previous INTEROPs, more than 20 vendors are expected to participate in this year's *Open Systems Interconnection* (OSI) demonstration. Over the years, the OSI showcase has evolved from an interoperability demonstration to an impressive forum, where vendors show worldwide connectivity on the exhibition floor. As was the case last year, the demonstration this year will involve computer systems located overseas in both Asia and Europe. This demonstration will focus on OSI solutions for the day-to-day business transactions in a multivendor, open systems environment. The scenarios include financial planning, direct mailings to a specific subset of a customer base, and the processing of product change requests. The application protocols used to perform these tasks are X.400 MHS (electronic mail), X.500 Directory Services, and FTAM (file transfer).

OSPF

Vendors will be using *Open Shortest Path First* (OSPF), the latest in interior routing protocol technology, to provide communication across the INTEROP 91 Fall Shownet. OSPF, the next generation routing standard for TCP/IP environments was developed to provide greater efficiency and control over how traffic is routed in large, multivendor networks. By supporting techniques such as least-cost and multipath routing, OSPF optimizes the use of network bandwidth and resources, both critical items in a growing internetwork. Visit this demonstration to see how OSPF can provide innovative solutions for your connectivity problems.

SMDS

"SMDS is what we need to get the true value out of data networks..." (Dan Lynch, *Communications Week*, 15 Oct. 90)—and the plan for the INTEROP 91 Fall *Switched Multimegabit Data Service* (SMDS) demonstration is to continue on this trajectory. SMDS has become a market reality through user interest, vendor announcements, new product capabilities, and service provider activities. SMDS at DS1 and DS3 speeds and the implementation of the IP-over-SMDS Internet standard, RFC 1209, will be demonstrated. Inclusion of SNMP-based Network Management capabilities is planned. A number of vendors' products will be interoperating on the show floor demo in a variety of applications.

SNMP

In the past few years, the *Simple Network Management Protocol* (SNMP) demonstration has succeeded in raising consumer awareness of the protocol and has illustrated the pervasiveness of SNMP implementations in our industry. We have honestly claimed that the products interoperate. This year, we will improve the value of the demonstration by effectively and precisely showing how SNMP products can be combined to provide manageable networks. Through this demonstration, the public will be educated on what SNMP does and how to use it, and will see how these offerings can provide manageable networks using SNMP. Using the INTEROP 91 Fall Shownet as the backbone of the demonstration, the SNMP booth will be the vehicle to communicate the message: "SNMP works for you."

Token Ring

The *IEEE 802.5 Token Ring* standards is a preferred technology for LANs due to its deterministic characteristics, high bandwidth, robustness, and built-in network management capabilities. Token Ring networks operate at 4 or 16 Mbps data rates over a variety of popular networking media, such as Shielded Twisted Pair (STP), Unshielded Twisted Pair (UTP), and Fiber Optic links. The INTEROP Token Ring Solution Showcase will illustrate Token Ring interoperability under "real world" conditions. High performance networking applications will be displayed on workstations within the booth. For the first time, INTEROP will be including Token Ring as part of its show network.

Book Reviews

Cyberpunk

Katie Hafner and John Markoff, *Cyberpunk—Outlaws and Hackers on the Computer Frontier,* Simon and Schuster, (New York: 1991), ISBN 0-671-68322-5, 368 pages.

Cyberpunk has the kind of cover that makes a technocrat cringe. Clifford Stoll of *The Cuckoo's Egg* tells us this is "An astonishing story," the subtitle tells us we will learn about outlaws and hackers on the computer frontier. Open the cover, however, and the sensationalism disappears.

The authors of this book have impeccable credentials for reporting on the computer underground. Katie Hafner has worked for *Business Week* and John Markoff is a well-respected, well-known reporter for the *New York Times.* The two set out to tell us about the "social consequences of computer networks and the communities that have grown up around them."

Three stories

Cyberpunk tells three stories. It starts with the tale of Kevin Mitnick, a notorious phone "phreak" who plagued computer administrators in Southern California for many years. Next, the book describes "Pengo," a member of the Chaos computer club which Clifford Stoll made famous. Lastly, the book describes Robert Morris, author of the *worm* that ate the Internet.

Mitnick

Mitnick's story is truly a scary one. In a well-balanced, well-researched account, we learn much of the background of Mitnick and his maladjusted friends as they begin wreaking havoc in the phone company. As the phone company migrated towards computer-based control systems, Mitnick learned more and more about computer systems.

What is fascinating about this story is the extent to which *people* are the security hole in the phone system. Mitnick repeatedly used his social engineering skills to call up operators and get system passwords. Once on the systems, however, Mitnick demonstrated formidable hacking skills.

Perhaps the most interesting hack by Mitnick is his theft of the source code for the VMS operating system, directly off the development VAX Cluster in New Hampshire. Even more interesting is the fact that DEC security were repeatedly informed of this theft and stood idly by. Only when Brian Reid of the Western Research Laboratory was notified did DEC finally swing into action.

Pengo

The second story is about Pengo. While Clifford Stoll starts with himself and tells of his own exploits, *Cyberpunk* starts in Germany and tells us about how this all came to be. Stoll finally tracked down Markus Hess, one of the more benign members of the Chaos computer club. Pengo (his handle comes from a popular video game) was a more sophisticated hacker than Hess and was more deeply involved in cracking activities.

Through another member of the Chaos club, a cooperative arrangement was set up with the KGB through contacts in East Germany. Everybody thought they would get rich off selling software to the Soviets, but reality was a little more mundane. Real software was hard to get, so the members of Project Equalizer started selling public domain software to the Soviets. Being technically unsophisticated, the Soviets were happy to shell out moderate amounts of money. Selling software to the Soviets was, rhetoric to the contrary, no ideological crusade.

Instead, the members were trying to finance their own habits. One wanted to buy hash, another wanted to buy fancy meals. Pengo simply wanted to upgrade his computer equipment from a Commodore to a VAX, a nice little terminal from which to base cracking activities.

Robert Morris

The third story in *Cyberpunk* is that of Robert T. Morris. *Cyberpunk* gives us valuable background information about Morris and his father, a noted security expert (Morris senior helped write the password encryption mechanism in UNIX). Morris junior is no technical slouch, either: while a sixteen-year old intern at Bell Labs he developed a new version of *uucp*.

The description of the Morris worm is notable for its technically accurate, thoughtful description of law and ethics in the Internet. We see that Morris was simply an outgrowth of a culture that encouraged such work (although, granted, it encouraged less sloppy instances of such work). Hafner and Markoff help defuse some of the viciousness behind the Morris lynching in the computer community by showing us how this episode came to be.

Ranks with the classics

Throughout the book, *Cyberpunk* takes a balanced and remarkably accurate view. The authors are technically sophisticated and objective, making *Cyberpunk* a significant contribution to the field, in addition to being a fascinating story. *Cyberpunk* ranks with classics like Tracy Kidder's *Soul of a New Machine* in bringing to life the arcane underground of computer networks.

Devouring Fungus

Karla Jennings, *The Devouring Fungus: Tales of the Computer Age*, Norton (New York: 1990), 237 pages.

The Devouring Fungus is one of the funniest books about computers this reviewer has read. In a well-written compilation of jokes, stories, and other materials from a wealth of sources, Karla Jennings has created an extremely entertaining look at our industry.

Did you know how many IBM technicians it takes to change a tire? *The Devouring Fungus* tells us: it takes one to jack up the car and another one to swap out the tires until they find the flat. Or, did you hear about the computer scientist who was asked what he did for a living? "I teach UNIX" he said. "Oh, that's great," was the reply, "What do you teach them?"

Amaze your friends and bore your spouse with *The Devouring Fungus*.

—*Carl Malamud*

Write to *ConneXions*!

Have a question about your subscription? Suggestions for topics? Want to write an article? A letter to the Editor? Have a question for an author? Want to enquire about back issues? (there are now more than fifty to choose from; ask for our free 1987–1991 index sheets). We want to hear from you. Simply write, call, fax, or e-mail to:

ConneXions—The Interoperability Report
480 San Antonio Road, Suite 100
Mountain View, CA 94040–1219
USA
Phone: +1 415-941-3399 or 1-800-INTEROP (Toll-free in the USA)
Fax: +1 415-949-1779
E-mail: `connexions@interop.com`

CONNEXIONS
480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

ADDRESS CORRECTION
REQUESTED